

EMC[®] VNX[™]

Open Systems Configuration Guide

P/N 300-013-444

REV 04

Copyright © 2011- 2014 EMC Corporation. All rights reserved. Published in the USA.

Published February, 2014

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

The information in this publication is provided as is. EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose. Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to the technical documentation and advisories section on the EMC online support website.

CONTENTS

Chapter 1	VNX Hardware	
	Introduction	12
	Hardware components.....	15
	VNX configurations by model	30
	VNX Gateway systems.....	80
	VNX Hardware Upgrades	89
	VNX for Block to Unified upgrades	90
	VNX Data in Place Conversions	91
 Chapter 2	 VNX General SW Configuration Rules	
	VNX series MirrorView ports.....	94
	Host connection to VNX series storage systems	95
	Fibre Channel storage area network (SAN) configuration rules.....	95
	VNX storage-system port speed rules.....	98
	Mixing Fibre Channel configurations	101
	Mixing Fibre Channel, FCoE, and iSCSI configurations	105
	iSCSI network configuration rules	110
	Virtual provisioning for VNX series storage systems	119
	FAST tiering for VNX series storage systems	124
	Disk power savings.....	128
	Unisphere and domain connections.....	128
	Bootimg from SAN guidelines	131
	Replicator V2 Incremental Attach	138
	Failover configuration rules.....	140
	Non-disruptive software installation (NDU) configuration rules	146
	Reserved LUN pool	147
	MirrorView/A and MirrorView/S configuration rules	147
	SAN Copy configuration rules.....	150
	LUN usage with replication applications	154
	VNX Snapshots.....	155
	VAAI support for NFS.....	157
	VASA support for VMware	159
	VAAI support for block	160
 Chapter 3	 VNX R31-R32 SW Configuration Rules	
	Fibre Channel storage area network (SAN) configuration rules.....	164

	iSCSI network configuration rules	164
	Read/write cache configurations.....	165
	FAST Cache for VNX storage systems	168
	Compression for VNX series storage systems	169
	Storage-system parameters for VNX series storage systems	170
	Server software rules	175
	SnapView configuration rules	178
	MirrorView/A and MirrorView/S limits.....	184
	SAN Copy configuration limits.....	185
	VNX RecoverPoint write Splitter usage.....	186
Chapter 4	VNX R33 SW Configuration Rules	
	Fibre Channel storage area network (SAN) configuration rules.....	190
	iSCSI network configuration rules	191
	FAST Cache for VNX storage systems	195
	Storage-system parameters for VNX series storage systems	196
	SnapView configuration rules	201
	VNX Snapshots.....	205
	MirrorView/A and MirrorView/S	206
	SAN Copy configuration limits.....	207
	Multi-Core Cache (MCC)	207
	Multi-Core RAID (MCR)	208
	Hot Spare Policy	208
	Symmetric Active/Active.....	209
	Deduplication and compression	209
Chapter 5	Rail Kits for VNX Family Systems	
	Rail kits	214
Chapter 6	VNX and Fibre Channel Switches	
	Fibre Channel switch zoning	218
	Fibre Channel and FCoE switch topology rules.....	225
Chapter 7	VNX Cable Types	
	VNX cable types.....	228
Index		

FIGURES

	Title	Page
1	VNX5100, VNX5300, and VNX5500 DPE, front view.....	16
2	VNX5100, VNX5300, and VNX5500 DPE, rear view.....	17
3	VNX5200/VNX5400/VNX5600/VNX5800/VNX7600 DPE, front view.....	18
4	VNX5200/VNX5400/VNX5600/VNX5800/VNX7600 DPE, rear view.....	18
5	VNX5700/VNX7500 SPE front and rear views.....	19
6	VNX8000 SPE front view	20
7	VNX8000 SPE rear view.....	20
8	2U DAE front and rear views	21
9	3U DAE front and rear views	22
10	4U DAE front and rear views	23
11	1U SPS front and rear views	25
12	2U SPS front and rear views	25
13	2U Li Ion SPS front and rear views.....	26
14	Control Station front and rear views	27
15	Control Station front and rear views	27
16	DME front view.....	28
17	DME rear view.....	28
18	VNX5100 rear view	31
19	VNX5200 Block configuration	33
20	VNX5200 File/Unified configuration, rear view	34
21	VNX5300 Block configuration	37
22	VNX5300 File/Unified configuration, rear view	38
23	VNX5400 Block configuration	41
24	VNX5400 File/Unified configuration, rear view	42
25	VNX5500 Block configuration, rear view	45
26	VNX5500 File/Unified hardware configuration, rear view.....	46
27	VNX5600 Block configuration	49
28	VNX5600 File/Unified configuration, rear view	50
29	VNX5700 Block configuration, rear view	53
30	VNX5700 File/Unified configuration, rear view	54
31	VNX5800 Block configuration	58
32	VNX5800 File/Unified configuration, rear view	59
33	VNX7500 File/Unified configuration, rear view	62
34	VNX7600 Block configuration	70
35	VNX7600 File/Unified configuration, rear view	71
36	VNX8000 Block configuration	74
37	VNX8000 Unified configuration.....	75
38	VG2 rear view	81

39	VG8 rear view	83
40	VG10 rear view	85
41	VG50 rear view	87
42	VNX over FCoE gateway configuration	99
43	FCoE ports	101
44	Server and storage system in both Fibre Channel SAN and direct attach configurations	103
45	Server and storage system in multiple FCoE switch configurations	104
46	Server connected to both an iSCSI and Fibre Channel storage system over a network	106
47	Server and storage system in both iSCSI network and direct attach configurations	108
48	VNX storage system with multiple connections	109
49	Direct connections between server and storage system	141
50	SAN configurations (FC switch connections between server and storage system) .	142
51	FCoE switch connections between server and storage system	143
52	Network configurations (iSCSI network connections between server and storage system)	144
53	VAAI High Level Architecture Diagram	159
54	Example of a Windows cluster in a SnapView configuration with tape backup	183
55	Single-Initiator zoning with multiple HBA-port servers.....	219
56	Single-initiator zoning with dual HBA-port servers.....	221
57	Single-initiator zoning with dual HBA-port and single HBA-port servers connected to a VNX series system	224

TABLES

	Title	Page
1	VNX1 storage configuration options.....	13
2	VNX2 storage configuration options.....	13
3	VNX I/O modules used in storage processors.....	29
4	VNX I/O modules used in data movers.....	29
5	VNX5100 storage processor configuration	32
6	VNX5200 storage processor configuration	35
7	VNX5200 Data Mover configuration	36
8	VNX5300 storage processor configuration	39
9	VNX5300 Data Mover configuration	40
10	VNX5400 storage processor configuration	43
11	VNX5400 Data Mover configuration	44
12	VNX5500 storage processor configuration	47
13	VNX5500 Data Mover configuration	48
14	VNX5600 storage processor configuration	51
15	VNX5600 Data Mover configuration	52
16	VNX5700 storage processor configuration	55
17	VNX5700 Data Mover configuration	57
18	VNX5800 storage processor configuration	60
19	VNX5800 Data Mover configuration	61
20	VNX7500 storage processor configuration	63
21	VNX7500 Data Mover configuration	69
22	VNX7600 storage processor configuration	72
23	VNX7600 Data Mover configuration	73
24	Back end connections for a base VNX8000 8 loop system.....	76
25	Back end loops for VNX8000 system upgraded in the field.....	76
26	VNX8000 Back end connections for 16 loop system.....	77
27	VNX8000 storage processor configuration	78
28	Front End I/O Module Sequence	79
29	VNX8000 Data Mover configuration	80
30	VG2 Data Mover configuration	82
31	VG8 Data Mover configuration	84
32	VG10 Data Mover configuration	86
33	VG50 Data Mover configuration	88
1	VNX Data in Place conversions.....	92
34	MirrorView ports for factory systems	94
35	Initiator records for Data Movers	98
36	Storage processor port speeds	98
37	FC/FCoE front end ports	101

38	Number of paths supported from a server	115
39	Number of paths supported from a server running Windows Server 2003 and Windows 2000	117
40	Virtual provisioning and traditional provisioning.....	121
41	Initial tier settings for LUNs in a pool (FAST enabler not installed)	125
42	FAST Policies for LUNs in a pool (FAST enabler installed).....	126
43	Interaction between initial tier placement settings and FAST policies.....	126
44	Unisphere feature support	129
45	Maximum boot servers per pool or RAID group consisting of n devices.....	132
46	Number of reserved LUNs	147
47	Server operating systems for storage systems in SAN Copy.....	151
48	LUN usage with replication applications	154
49	Initiator records for Fibre Channel	164
50	Initiator records for iSCSI	165
51	Maximum cache allocations	166
52	FAST Cache and disk configurations.....	168
53	Storage system parameters	170
54	Server operating systems.....	175
55	SnapView Limits	178
56	SnapView supported limits for VNX series storage systems.....	179
57	Maximum SnapView snapshot limits for VNX series storage systems.....	181
58	MirrorView/A parameter limits for VNX series storage systems.....	184
59	MirrorView/S parameter limits for VNX series storage systems.....	185
60	SAN Copy parameter limits for VNX series systems.....	185
61	Interoperability with MirrorView/A, and MirrorView/S, SAN Copy, and SnapView .	186
62	RecoverPoint limits	187
63	Initiator records for Fibre Channel	190
64	Initiator records for iSCSI	191
65	Number of paths supported from a server.....	192
66	Number of paths supported from a server running Windows Server 2003 and Windows 2000	194
67	FAST Cache and disk configurations.....	195
68	Storage system parameters	196
69	SnapView supported limits for VNX series storage systems.....	202
70	Maximum SnapView snapshot limits for VNX series storage systems.....	204
71	VNX Snapshot parameters	205
72	MirrorView/A and MirrorView/S parameters.....	206
73	SAN Copy limits	207
74	VNX5100 and VNX5300 rail kits.....	214
75	VNX5500, VNX5700, and VNX7500 rail kits	214
76	VNX5200, VNX5400, VNX5600, VNX5800, VNX7600 and VNX8000 rail kits .	215
77	VNX5100 and VNX5300 TELCO rail kits	215

78	VNX5500, VNX5700, and VNX7500 TELCO rail kits.....	215
79	VNX5200, VNX5400, VNX5600, VNX5800, and VNX7600 TELCO rail kits	216
80	Zone contents for storage-system Mirror Ports	222
81	Zone contents for server HBAs	222
82	Zone contents.....	225
83	Optical OM3 cables for 8 Gb, 4 Gb, or 2 Gb operation	228
84	Active TwinAX cable for FCoE or 10Gb iSCSI operation.....	229
85	VNX5100, VNX5300, VNX5500 DPE to DAE SAS cables	230
86	DPE/SPE to DAE SAS cables	230
87	SPE to DAE SAS cables.....	230
88	DAE to DAE SAS cables for VNX1 series	231
89	DAE to DAE SAS cables for VNX2 series	231

CHAPTER 1

VNX Hardware

This section provides a high-level overview of the VNX series hardware. It includes both individual components and complete storage systems. Topics include:

- ◆ Introduction 12
- ◆ Hardware components 15
- ◆ VNX configurations by model..... 30
- ◆ VNX Hardware Upgrades 89
- ◆ VNX for Block to Unified upgrades 90
- ◆ VNX Data in Place Conversions 91

Introduction

The VNX family of storage systems consists of models that fall into two categories: the Unified Family, and the VNX Gateways.

Unified family

The Unified Family consists of the following models:

- ◆ VNX5100
- ◆ VNX5200
- ◆ VNX5300
- ◆ VNX5400
- ◆ VNX5500
- ◆ VNX5600
- ◆ VNX5700
- ◆ VNX5800
- ◆ VNX7500
- ◆ VNX7600
- ◆ VNX8000

The Unified Family models provide the following types of storage capabilities:

- ◆ **Block:** VNX for Block storage systems provide storage capability to Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), and internet small computer system interface (iSCSI) hosts. Block storage systems consist of storage processors, disk drives, and (in some models) standby power supplies.

Note: The VNX5100 does not support block storage over FCoE or iSCSI.

- ◆ **File:** VNX for File storage systems provide NFS and CIFS storage capabilities for network clients over copper and optical Ethernet connections. File storage systems consist of Control Stations, Data Movers, storage processors, disk drives, and (in some models) standby power supplies.

Note: The Data Movers and storage processors are connected to each other with FC or FCoE connections.

- ◆ **Unified:** VNX Unified storage systems provide both Block and File storage capabilities. Unified storage systems consist of Control Stations, Data Movers, storage processors, disk drives, and (in some models) standby power supplies.

[Table 1](#) and [Table 2](#) list the available storage configurations for the VNX Unified Family.

Table 1 VNX1 storage configuration options

Storage system	Block	File	Unified
VNX5100	Y	N	N
VNX5300	Y	Y	Y
VNX5500	Y	Y	Y
VNX5700	Y	Y	Y
VNX7500	Y	Y	Y

VNX1 hardware runs Block OE versions 5.31/5.32 and File OE versions 7.0/7.1.

Table 2 VNX2 storage configuration options

Storage system	Block	File	Unified
VNX5200	Y	Y	Y
VNX5400	Y	Y	Y
VNX5600	Y	Y	Y
VNX5800	Y	Y	Y
VNX7600	Y	Y	Y
VNX8000	Y	Y	Y

VNX2 hardware runs Block OE version 5.33 and File OE version 8.1.

[“Hardware components” on page 15](#) provides more information about the VNX hardware components.

VNX gateways

There are four VNX gateway models. They are:

VNX1:

- ◆ VG2

- ◆ VG8

VNX2:

- ◆ VG10

- ◆ VG50

VNX gateway systems provide CIFS and NFS access to network clients over copper and optical Ethernet connections, but do not provide any storage capability. Gateways require a separate storage array to provide disk space and to store all required boot information. VNX VG2 and VG8 gateways can connect to one or more EMC CLARiiON systems, EMC Symmetrix systems, or EMC VNX for Block systems. VNX VG10 and VG50 can connect to one or more EMC Symmetrix systems. Gateway systems consist of Control Stations and Data Movers.

[“Hardware components” on page 15](#) provides more information about the VNX hardware components.

Specialized variants

There are two specialized configuration options available for VNX storage systems:

- ◆ **DC power:** DC-powered systems are intended for use in environments that require redundant, highly-available power sources such as a telecommunications central office. DC-powered VNX systems use the same basic hardware components, but with DC power supplies instead of the standard AC power supplies. The DC power option is available on the VNX5100 and the Block, File, and Unified versions of the VNX5200, VNX5300, VNX5400, VNX5500, and VNX5600.
- ◆ **Dense storage:** The dense storage option is available on the Block, File, and Unified versions of the VNX5500, VNX5700, and VNX7500 models. Dense configurations are also available for all configurations of the VNX5400, VNX5600, VNX5800, VNX7600, and VNX8000. The dense storage option decreases the cabinet space required to connect the maximum number of disks to the VNX system, but it does not increase the maximum number of disks supported. The dense storage option uses a 4U 60-drive disk array enclosure (DAE), and a 40U dense cabinet that is deeper and requires more front access space than a standard 40U cabinet.

Hardware components

This section provides an overview of the VNX storage processors, disk array enclosures, standby power supplies, Control Stations, Data Movers, and IO modules in the VNX hardware.

Disk processor enclosure (DPE) and storage processor enclosure (SPE)

All VNX Unified Family systems, regardless of whether they are Block, File, or Unified configurations, have storage hardware that consists of storage processors and disk drives.

The storage hardware for each VNX system uses a disk processor enclosure (DPE) or a storage processor enclosure (SPE). Although each enclosure performs a similar function, there are distinct differences in the physical layouts of the enclosures.


Each DPE or SPE contains dual storage processors.

DPE


A DPE combines the first shelf of disks, including the vault drives, and the storage processors into a single enclosure. The disks are mounted in the front of the enclosure, and the storage processor components are in the back of the enclosure.

Figure 1 shows the front of the DPE used in the VNX5100, VNX5300, and VNX5500, and Figure 2 on page 17 shows the rear of that DPE.

3U, 25 2.5" drive Disk processor enclosure front




3U, 15 3.5" drive Disk processor enclosure front



VNX-000140

Figure 1 VNX5100, VNX5300, and VNX5500 DPE, front view

Disk processor enclosure rear



VNX-000141

Figure 2 VNX5100, VNX5300, and VNX5500 DPE, rear view

In these DPE-mounted storage processors, the mezzanine ports provide the connections for the minimum system configuration. There are also 2 slots to install optional I/O modules if additional connectivity is required.

The following ports are located on the DPE mezzanine:

- ◆ Two SAS ports for BE bus connections (ports 0 and 1)
- ◆ Four Fibre-Channel ports for blade or host connections (ports 2-5)

Note: Although there are four FC ports and two expansion slots for I/O modules per SP on the DPE, the supported usage of those ports and slots varies between models. Refer to [“VNX5100” on page 31](#), [“VNX5300” on page 36](#), or [“VNX5500” on page 44](#) for specific information about each VNX system that uses a DPE.

[Figure 3 on page 18](#) shows the front view of the DPE used in the VNX5200, VNX5400, VNX5600, VNX5800, and VNX7600. [Figure 4](#) shows the rear view of that DPE. This DPE utilizes only 2.5” disks, has 2 SAS ports (labeled 0, 1), and has no mezzanine card. All other ports are available using I/O modules.

Disk-processor enclosure front




Figure 3 VNX5200/VNX5400/VNX5600/VNX5800/VNX7600 DPE, front view

Disk-processor enclosure Rear




Figure 4 VNX5200/VNX5400/VNX5600/VNX5800/VNX7600 DPE, rear view

Note: The VNX5200 DPE has “Do Not Use” covers over slots 0 and 1 of each SP. The VNX5400 DPE uses the “Do Not Use” module cover for slot 0. All slots are supported in other models.

SPE

An SPE contains only the storage processor components. It does not contain any disk drives. The enclosure fans and power supplies are located in the front of the enclosure, and the ports are located in the back of the enclosure.

SPE-mounted storage processors do not have any mezzanine ports. All of the ports for storage connections are on I/O modules. The number of I/O modules that are included to support the minimum system configuration, and the number of optional I/O modules vary between the individual VNX models. Most systems will have a SAS module in slot 0 (as shown in Figure 5) for connections to DAEs. Those used in File/Unified configurations will also have a Fibre Channel module in slot 4 (for connecting to the data movers).

Figure 5 shows the front and rear of the SPE used in the VNX5700 and VNX7500.




Figure 5 VNX5700/VNX7500 SPE front and rear views

SPE Front




Figure 6 VNX8000 SPE front view

Figure 6 and **Figure 7** show the SPE used in the VNX8000. It contains 11 slots/SP for I/O modules, numbered 0-10.



Figure 7 VNX8000 SPE rear view


Disk array enclosure (DAE)

Disk array enclosures are hardware components that house disk drives for VNX Block, File, and Unified systems. There are three different types of DAEs.

2U DAE (DAE5S)

Figure 8 shows the front and rear of the 2U 25 2.5" disk DAE.

2U, 25 2.5" drive DAE Front



Rear

VNX-000227

Figure 8 2U DAE front and rear views

3U DAE (DAE6S)

Figure 9 shows the front and rear of the 3U 15 3.5" disk DAE.




Figure 9 3U DAE front and rear views

4U DAE (DAE7S)

The VNX Disk Array Enclosure (DAE) containing 60 3.5" 6G SAS disk drives in a 40U deep rack cabinet is used in VNX5500, VNX5700, and VNX7500 dense systems. It can also be used in VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, and VNX8000 systems.

Note: The 60 drive DAE cannot be installed in a traditional non-deep 40U rack.

Figure 10 shows the front and rear of the 4U DAE. The 4U DAE can hold up to 60 disks. In this DAE, the drives are loaded from the top.

IMPORTANT

The 4U DAE requires a deep rack and is not customer-installable.

4U 60 drive Disk-array enclosure




Figure 10 4U DAE front and rear views

The VNX 4U Disk Array Enclosure (DAE) containing 60 disk drives in a 40U deep rack cabinet is used in VNX5200, VNX5400, VNX5500, VNX5600, VNX5700, VNX5800, VNX7500, VNX7600, and VNX8000 dense configurations.

Additional information on dense configurations is available in model-specific dense configuration install addenda and other documentation available at:

- ◆ <http://mydocs.emc.com/VNX/>
- ◆ the VNX Procedure Generator

DAE Configuration Rules

- ◆ In general, the 3U DAE6S and the 2U DAE5S can be combined in any sequence.
- ◆ Each system has a maximum disk drive capacity enforced. The operating system and USM will not allow addition of DAEs that would provide more disks that the system can support.
- ◆ In general, one BE SAS loop supports up to 10 DAEs.
- ◆ The design of the storage array's DAEs depends upon numerous issues, including FAST cache usage; the design and usage of tiered storage; the importance of disk density and power options; etc. See [“FAST tiering for VNX series storage systems” on page 124](#), [“Disk power savings” on page 128](#), and other resources in [“VNX General SW Configuration Rules”](#).
- ◆ The 3U DAE6S and the 2U DAE5S fit into the standard 40U EMC cabinet. The 4U DAE7S requires a deep cabinet.
- ◆ The DAE6S can be intermixed with the DAE7S. Check the *VNX Disk and OE Support Matrix* on the EMC Online Support website for allowable disk drive configurations.
- ◆ In general, DAE5S cannot be used with 40U deep rack cabinets since drive capacities and speeds of the DAE are not configurable in the DAE7S. However, dense configurations using Flash drives for FAST cache may have those drives in a “Primary” DAE5S DAE to enable maximum FLASH performance.
- ◆ The spindle and capacity density of the 60 x 3.5” DAEs (60 drives/4U) are better than the 25 x 2.5” DAEs (25 drives/2U).

Standby power supply (SPS)

Standby power supplies are battery backups on the VNX Block, File, and Unified systems that enable the array time to shut down gracefully in the event of a loss of power. All VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500 systems in standard cabinets use a 1U SPS. VNX5500, VNX5700, and VNX7500 systems use a 2U SPS when a 4U 60-drive DAE serves as DAE 0. The VNX5200, VNX5400, VNX5600, VNX5800, and VNX7600 do not require any SPS because they have a Battery on Bus unit mounted inside each SP. The VNX8000 utilizes two Li Ion 2U SPSs.

VNX5100 and VNX5300 systems can have one or two 1U SPSs. VNX5500, VNX5700, and VNX7500 systems have two 1U SPSs.

Figure 11 shows the front and rear views of the 1U SPS used with VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500 systems in standard cabinets..




Figure 11 1U SPS front and rear views

Figure 12 shows the front and rear views of the 2U SPS used with VNX5500, VNX5700, and VNX7500 dense configurations utilizing the 4U DAE as the vault DAE.




Figure 12 2U SPS front and rear views

Figure 13 shows the front and rear views of the 2U Li Ion SPS used in the VNX8000. Each of those systems uses two of these SPSs.




Figure 13 2U Li Ion SPS front and rear views

Control Station (CS)

The Control Station is the hardware component that provides centralized management functions for VNX File, Unified, and gateway systems. The Control Station monitors the status of the Data Movers and storage processors over an internal network, and connects to the external network for management purposes.

VNX for File/Unified and gateway systems can have one or two Control Stations. For systems with two Control Stations, the first CS [designated as CS 0] is the primary, and the second [CS 1] serves as a standby. A crossover cable connects the two Control Stations, and allows the standby CS to take over if the primary fails.

[Figure 14](#) shows the front and rear views of the Control Station used in the VNX5300, VNX5500, VNX5700, VNX7500, VG2, and VG8.





Figure 14 Control Station front and rear views

Figure 15 shows the front and rear views of the Control Station used in the VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, VNX8000, VG10 and VG50.

Control Station Front



Control Station Rear




Figure 15 Control Station front and rear views


Data Mover enclosure (DME)

Data Movers are the hardware components for VNX File and Unified configurations that provide the interface between the storage array and users accessing the VNX over the external network. A DME can contain one or two Data Movers.

The enclosure fans and power supplies are located in the front of the enclosure, and the ports are located in the back of the enclosure.

Data Movers do not have any mezzanine ports. All of the ports for storage and network connections are on I/O modules. The number of I/O modules that are included to support the minimum system configuration, and the number of optional I/O modules vary between the individual VNX models.

Figure 16 shows the front of the DME, and Figure 17 shows the rear of the DME.



Blade enclosure front

VNX-000138

Figure 16 DME front view




Figure 17 DME rear view

VNX-000139

I/O modules

Table 3 shows the VNX I/O modules and lists whether the I/O modules can be used in the storage processors, blades, or both.

Table 3 VNX I/O modules used in storage processors

I/O module type	Supported in VNX OE ^a	Connector type
4-port 8 Gb Fibre Channel I/O module	R31, R32, R33	SFP+
4-port 1 Gb copper iSCSI I/O module	R31, R32, R33	RJ45
2-port 10 Gb optical or active TwinAx iSCSI I/O module	R31, R32, R33	SFP+
2-port 10 Gb optical or active TwinAx FCoE I/O module	R31, R32, R33	SFP+
4-port 6 Gb SAS I/O module	R31, R32, R33	Mini-SAS HD
2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	R31, R32, R33	RJ45

a. While these I/O modules are supported in these releases, there is no software upgrade path from R32 to R33. Effectively, R33 starts a new and different software (and hardware) path.

Table 4 lists the VNX I/O modules used in data movers.

Table 4 VNX I/O modules used in data movers

I/O module type	Supported in VNX OE ^a	Connector type
4-port 8 Gb Fibre Channel I/O module	7.0, 7.1, 8.1	SFP+
4-port 1 Gb copper Ethernet I/O module	7.0, 7.1, 8.1	RJ45
2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	7.0, 7.1	SFP+, RJ45
2-port 10 Gb optical Ethernet I/O module	7.0, 7.1, 8.1	SFP+

Table 4 VNX I/O modules used in data movers

I/O module type	Supported in VNX OE ^a	Connector type
2-port 10 Gb optical Ethernet I/O module	7.1	SFP+
2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	7.1, 8.1	RJ45
2-port 10 Gb optical or active TwinAx FCoE I/O module	7.0, 7.1, 8.1	SFP+

a. While these I/O modules are supported in these releases, there is no software upgrade path from 7.1 to 8.1. Effectively, 8.1 starts a new and different software (and hardware) path.

I/O module port identification

Unisphere assigns consecutive logical port identifiers (A-0, A-1, A-2, A-3... for SP A FE ports and B-0, B-1, B-2, B-3... for SP B FE ports) starting with the lowest numbered FE port that has been initialized on the I/O module in the lowest numbered SP slot. The ports on all I/O modules that ship with a storage system are set and persisted at the time the system is initialized. Note that the initialization process writes the port configuration information to persistent memory.

For more information on the I/O modules available for use in the storage processors or data movers, see the Hardware Information Guide for each specific model.

VNX configurations by model

This section breaks down the hardware for each VNX model.

For a much more detailed treatment of each VNX model's hardware, see the Hardware Information Guide for that specific model. These documents are available online at:

<https://mydocs.emc.com/VNX/requestMyDoc.jsp>

or

https://support.emc.com/products/12781_VNX-Series

VNX5100

The VNX5100 only supports block storage. It does not come in File or Unified configurations. It cannot be upgraded to a Unified configuration. It is also available in a DC-powered model.




Figure 18 VNX5100 rear view

The VNX5100 SPs do not support I/O modules for additional connections.

Note: The VNX5100 DPE provides four FC ports per SP, but only two ports come with the required SFPs installed. Additional SFPs must be ordered to use the two remaining FC ports.

Table 5 shows the SP mezzanine ports for the VNX5100.

Table 5 VNX5100 storage processor configuration

Configuration	Mezzanine port type and quantity
Block	Two SAS ports for BE bus connections
	4 Fibre Channel ports Two Fibre Channel ports with SFP modules for host connections + 2 Two Fibre Channel ports without SFP modules.
	Note: Additional SFPs must be ordered to use the two remaining FC ports.

VNX5200

This section describes the hardware for the VNX5200. The VNX5200 is available in AC and DC-powered Block, File, and Unified configurations. Figure 19 illustrates a VNX5200 Block configuration.

Note: Slots 0 and 1 are covered in each SP as no I/O modules are supported in those slots.

The VNX5200 uses a DPE which includes a Battery on Bus; as a result, no SPS is necessary. The DPE contains the first 25 disks.




Figure 19 VNX5200 Block configuration

Figure 20 shows a VNX5200 File/Unified configuration rear view.




Figure 20 VNX5200 File/Unified configuration, rear view

Storage processor configuration

The VNX5200 SPs can be upgraded with additional I/O modules for additional FC, FCoE, or iSCSI connections.

For all configurations, there are 2 SAS ports (0, 1) for BE bus connections. The VNX5200 does not support additional SAS ports.

Table 10 shows the available I/O module configurations for the VNX5200.

Table 6 VNX5200 storage processor configuration

Configuration	Required I/O module		Optional I/O modules		
	I/O module type	Location	I/O module type	Maximum additional modules supported per SP	Location
Block	No requirement		4-port 8 Gb Fibre Channel I/O module	3 (Block) 3 (File/Unified)	Slots 2-4 are available if Block; slots 2-3 are available if File/Unified.
File/Unified	4-port 8 Gb Fibre Channel I/O module	Slot 4	4-port 10 GbE iSCSI I/O module	3	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	3 (Block) 3 (File/Unified)	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	3	
			4-port 1 GBase-T iSCSI I/O module	3	

Note: Both SPs must have the same modules in the same slots.

Data Mover configuration

Table 11 shows the Data Mover I/O module requirements and options for the VNX5200 File and Unified configurations.

Note: The VNX5200 DME supports I/O modules in 3 slots, slot numbers 0-2. Slots 3 and 4 have covers over them as they are not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 7 VNX5200 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1GbBase T for IP I/O module	2	Slot 1 Slot 2
			2-port 10 Gb (optical & active twinax) Ethernet I/O module	2	
			2-port 10 Gb Base-T IP I/O module	2	

VNX5300

The VNX5300 is available in Block, File, and Unified configurations. All of these configurations are also available in DC-powered systems. [Figure 21](#) illustrates a VNX5300 Block configuration.




Figure 21 VNX5300 Block configuration

Figure 22 shows a VNX5300 File/Unified configuration rear view.




Figure 22 VNX5300 File/Unified configuration, rear view

Storage processor configuration

The VNX5300 SPs can be upgraded with up to two I/O modules for additional FC, FCoE, or iSCSI connections.

For all configurations, there are 2 SAS ports (0, 1) for BE bus connections and 4 Fibre Channel ports (2, 3, 4, 5).

Although there are four FC ports on the mezzanine of each SP and four additional FC ports available to each SP by adding an optional I/O module, a maximum of two FC ports per SP are supported for blade connections in a File or Unified configuration. Two of the FC ports on the mezzanine card are reserved for the blade connections. Using those ports for host connections to a Block configuration ensures that the system cannot be upgraded to a Unified configuration.

Table 8 shows the available I/O module configurations for the VNX5300.

Table 8 VNX5300 storage processor configuration

Configuration	Optional I/O modules		
	I/O module type	Maximum additional modules supported per SP	Location
All	4-port 8 Gb Fibre Channel I/O module	1	Slot 0 Slot 1
	4-port 1 Gb copper iSCSI I/O module	1	
	2-port 10 Gb optical/-TwinAx FCoE I/O module	2	
	2-port 10 Gb optical/-TwinAx iSCSI I/O module	2	
	2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	1	

Note: Both SPs must have the same modules in the same slots.

Data Mover configuration

Table 9 shows the Data Mover I/O module requirements and options for the VNX5300 File and Unified configurations.

Note: The VNX5300 DME supports I/O modules in 3 slots, slot numbers 0-2. Slots 3 and 4 have covers over them as they are not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 9 VNX5300 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	2	Slot 1 Slot 2
			2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	2	
			2-port 10 Gb optical Ethernet I/O module	2	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	2	
			4-port Fibre Channel I/O module	2	

Figure 17 on page 28 shows the location of the DME I/O module slots.

VNX5400

This section describes the hardware for the VNX5400. The VNX5400 is available in AC- and DC-powered Block, File, and Unified configurations. Figure 23 illustrates a VNX5400 Block configuration.

Note: Slot 0 is covered in each SP as no I/O module is supported in that slot.

The VNX5400 uses a DPE which includes a Battery on Bus; as a result, no SPS is necessary. The DPE contains the first 25 disks.




Figure 23 VNX5400 Block configuration

Figure 24 shows a VNX5400 File/Unified configuration rear view.




Figure 24 VNX5400 File/Unified configuration, rear view

Storage processor configuration

The VNX5400 SPs can be upgraded with additional I/O modules for additional FC, FCoE, or iSCSI connections.

For all configurations, there are 2 SAS ports (0, 1) for BE bus connections. The VNX5400 does not support additional SAS ports.

Table 10 shows the available I/O module configurations for the VNX5400.

Table 10 VNX5400 storage processor configuration

Configuration	Required I/O module		Optional I/O modules		
	I/O module type	Location	I/O module type	Maximum additional modules supported per SP	Location
Block	No requirement		4-port 8 Gb Fibre Channel I/O module	4 (Block) 3 (File/Unified)	Slots 1-3 if File/Unified Slots 1-4 if Block
File/Unified	4-port 8 Gb Fibre Channel I/O module	Slot 4	4-port 10 GbE iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	4 (Block) 3 (File/Unified)	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	2	
			4-port 1 GBase-T iSCSI I/O module	2	

Note: Both SPs must have the same modules in the same slots.

Data Mover configuration

Table 11 shows the Data Mover I/O module requirements and options for the VNX5400 File and Unified configurations.

Note: The VNX5400 DME supports I/O modules in 3 slots, slot numbers 0-2. Slots 3 and 4 have covers over them as they are not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 11 VNX5400 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1GbBase T for IP I/O module	2	Slot 1 Slot 2
			2-port 10 Gb (optical & active twinax) Ethernet I/O module	2	
			2-port 10 Gb Base-T IP I/O module	2	

VNX5500

The VNX5500 is available in standard (AC--power), DC-powered, and in dense configurations. In addition, the VNX5500 has a high-performance configuration with an additional SAS module, providing up to 4 more back end loops (and additional bandwidth). It does not change the disk limit for the model.

For more information on specialized configurations, see:

- ◆ *VNX5500 Dense Cabinet Installation Addendum*
- ◆ *DC-Powered VNX Series Enclosures Installation and Operation Guide*

Figure 25 shows the VNX5500 Block configuration hardware, rear view.





Figure 25 VNX5500 Block configuration, rear view

Figure 26 shows a VNX5500 File/Unified configuration rear view.



VNX-000146

Figure 26 VNX5500 File/Unified hardware configuration, rear view

Storage processor configuration

The VNX5500 SPs can be upgraded with up to two I/O modules for additional FC, FCoE, or iSCSI connections.

For environments with high bandwidth requirements, one of the I/O module slots on each SP can be used for a SAS I/O module that provides four additional BE buses.

Table 12 shows the available I/O module configurations for the VNX5500 configurations.

Table 12 VNX5500 storage processor configuration

Configuration	Optional I/O module slots per SP		
	I/O module type	Maximum additional modules supported per SP	Location
All	4-port 8 Gb Fibre Channel I/O module	1	Slot 0 Slot 1
	4-port 1 Gb copper iSCSI I/O module	2	
	2-port 10 Gb optical/-TwinAx FCoE I/O module	2	
	2-port 10 Gb optical/-TwinAx iSCSI I/O module	2	
	2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	1	
	4-port 6 Gb SAS I/O module	1	

Note: Both SPs must have the same modules in the same slots.

Although there are four FC ports on the mezzanine of each SP and four additional FC ports available to each SP by adding an optional I/O module, a maximum of three FC ports per SP are supported for blade connections in a File or Unified configuration.

Data Mover configuration

Table 13 shows the Data Mover I/O module configurations for the VNX5500 File and Unified configurations.

Note: The VNX5500 DME supports I/O modules in 4 slots, slot numbers 0-3. Slot 4 has a cover over it as it is not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 13 VNX5500 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	3	Slot 1 Slot 2 Slot 3
			2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	3	
			2-port 10 Gb optical Ethernet I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	3	

VNX5600

This section describes the hardware for the VNX5600. The VNX5600 is available in Block, File, and Unified configurations. [Figure 27](#) illustrates a VNX5600 Block configuration.

The VNX5600 uses a DPE which includes a Battery on Bus; as a result, no SPS is necessary. The DPE contains the first 25 disks.




Figure 27 VNX5600 Block configuration

Figure 28 shows a VNX5600 File/Unified configuration rear view.




Figure 28 VNX5600 File/Unified configuration, rear view

Storage processor configuration

The VNX5600 SPs can be upgraded with additional I/O modules for additional FC, FCoE, SAS, or iSCSI connections.

For all configurations, there are 2 SAS ports (0, 1) for BE bus connections. The VNX5600 does support the addition of 4 ports using a SAS I/O module.

The VNX5600 SP supports up to 5 I/O modules in each SP.

Table 14 shows the available I/O module configurations for the VNX5600.

Table 14 VNX5600 storage processor configuration

Configuration	Required I/O module		Optional I/O modules		
	I/O module type	Location	I/O module type	Maximum additional modules supported per SP	Location
Block	No requirement		4-port 8 Gb Fibre Channel I/O module	5 (Block) 4 (File/Unified)	Slots 0-3 if File/Unified Slots 0-4 if Block
File/Unified	4-port 8 Gb Fibre Channel I/O module	Slot 4	4-port 10 GbE iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	5 (Block) 4 (File/Unified)	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	2	
			4-port 1 GBase-T iSCSI I/O module	2	
			4-port 4x lane/port 6 Gb SAS I/O module	1	Slot 0

Note: Both SPs must have the same modules in the same slots.

Data Mover configuration

Table 15 shows the Data Mover I/O module requirements and options for the VNX5600 File and Unified configurations.

Note: The VNX5600 DME supports I/O modules in 3 slots, slot numbers 0-2. Slots 3 and 4 have covers over them as they are not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 15 VNX5600 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1GbBase T for IP I/O module	2	Slot 1 Slot 2
			2-port 10 Gb (optical & active twinax) Ethernet I/O module	2	
			2-port 10 Gb Base-T IP I/O module	2	

VNX5700

The VNX5700 comes in Block, File, and Unified configurations. It is also available in a dense configuration utilizing the 4U DAE as the vault DAE in a dense cabinet.

For more information on specialized configurations, see:

- ◆ *VNX5700 Dense Cabinet Installation Addendum*

Figure 29 shows a VNX5700 Block configuration.




Figure 29 VNX5700 Block configuration, rear view

Figure 30 shows a VNX5700 File/Unified configuration rear view.




Figure 30 VNX5700 File/Unified configuration, rear view

Storage processor configuration

The VNX5700 uses an SPE. [Table 16](#) shows the SP I/O module configurations for the VNX5700 Block, File and Unified configurations.

Table 16 VNX5700 storage processor configuration (page 1 of 2)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
Block	One 4-port SAS I/O module for BE bus connections	Slot 0	4-port 8 Gb Fibre Channel I/O module	3	Slot 1 Slot 2 Slot 3 Slot 4
			4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	3	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module"	3	

Table 16 VNX5700 storage processor configuration (page 2 of 2)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
File and Unified	One 4-port SAS I/O module for BE bus connections	Slot 0	4-port 8 Gb Fibre Channel I/O module	2	Slot 1 Slot 2 Slot 3
	One 4-port Fibre Channel I/O module for blade connections	Slot 4	4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	3	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module"	3	

IMPORTANT

To upgrade a Block system to a File or Unified configuration after installation, slot 4 in each SP must be left open for a Fibre Channel I/O module.

Data Mover configuration

Table 17 shows the Data Mover I/O module configurations for the VNX5700 File and Unified configurations.

Table 17 VNX5700 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	3	Slot 1 Slot 2 Slot 3
			2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	3	
			2-port 10 Gb optical Ethernet I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	3	

VNX5800

This section describes the hardware for the VNX5800. The VNX5800 is available in Block, File, and Unified configurations. Figure 31 illustrates a VNX5800 Block configuration.

The VNX5800 uses a DPE which includes a Battery on Bus; as a result, no SPS is necessary. The DPE contains the first 25 disks.




Figure 31 VNX5800 Block configuration

Figure 32 shows a VNX5800 File/Unified configuration rear view.




Figure 32 VNX5800 File/Unified configuration, rear view

Storage processor configuration

The VNX5800 SPs can be upgraded with additional I/O modules for additional FC, SAS, FCoE, or iSCSI connections.

For all configurations, there are 2 SAS ports (0, 1) for BE bus connections. The VNX5800 does support the addition of 4 ports using a SAS I/O module.

The VNX5800 SP supports up to 5 I/O modules in each SP.

Table 18 shows the available I/O module configurations for the VNX5800.

Table 18 VNX5800 storage processor configuration

Configuration	Required I/O module		Optional I/O modules		
	I/O module type	Location	I/O module type	Maximum additional modules supported per SP	Location
Block	No requirement		4-port 8 Gb Fibre Channel I/O module	5 (Block) 4 (File/Unified)	Slots 0-3 if File/Unified Slots 0-4 if Block
File/Unified	4-port 8 Gb Fibre Channel I/O module	Slot 4	4-port 10 GbE iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	5 (Block) 4 (File/Unified)	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	2	
			4-port 1 GBase-T iSCSI I/O module	2	
			4-port 4x lane/port 6 Gb SAS I/O module	1	Slot 0

Note: Both SPs must have the same modules in the same slots.

Data Mover configuration

Table 19 shows the Data Mover I/O module requirements and options for the VNX5800 File and Unified configurations.

Note: The VNX5800 DME supports I/O modules in 4 slots, slot numbers 0-3. Slot 4 has a cover over it as it is not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 19 VNX5800 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1GbBase T for IP I/O module	3	Slot 1 Slot 2 Slot 3
			2-port 10 Gb (optical & active twinax) Ethernet I/O module	3	
			2-port 10 Gb Base-T IP I/O module	3	

VNX7500

The VNX7500 comes in Block, File, and Unified configurations. It is also available in a dense configuration utilizing the 4U DAE as the vault DAE in a dense cabinet.

For more information on specialized configurations, see:

- ◆ *VNX7500 Dense Cabinet Installation Addendum*

The VNX7500 Block configuration looks much like the VNX5700 Block configuration shown in [Figure 29 on page 53](#). [Figure 33](#) shows a VNX7500 File/Unified configuration rear view.




Figure 33 VNX7500 File/Unified configuration, rear view

Storage processor configuration

“SPE” on page 19 provides more information about the SPE hardware.

Table 20 shows the SP I/O module configurations for the VNX7500 Block, File, and Unified configurations.

Table 20 VNX7500 storage processor configuration (page 1 of 6)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
Block with four BE buses	One 4-port SAS I/O module for BE bus connections	Slot 0	4-port 8 Gb Fibre Channel I/O module	4	Slot 1 Slot 2 Slot 3 Slot 4
			4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	4	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module”	3	

Table 20 VNX7500 storage processor configuration (page 2 of 6)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
Block with eight BE buses	Two 4-port SAS I/O modules for BE bus connections	Slot 0 Slot 1	4-port 8 Gb Fibre Channel I/O module	3	Slot 2 Slot 3 Slot 4
			4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	3	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module"	3	

Table 20 VNX7500 storage processor configuration (page 3 of 6)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
File and Unified with four BE buses and two to four blades	One 4-port SAS I/O module for BE bus connections	Slot 0	4-port 8 Gb Fibre Channel I/O module	3	Slot 1 Slot 2 Slot 3
	One 4-port Fibre Channel I/O module for blade connections	Slot 4	4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	3	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	3	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module"	3	

Table 20 VNX7500 storage processor configuration (page 4 of 6)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
File and Unified with four BE buses and five to eight blades	One 4-port SAS I/O module for BE bus connections	Slot 0	4-port 8 Gb Fibre Channel I/O module	2	Slot 1 Slot 2
	Two 4-port Fibre Channel I/O modules for blade connections	Slot 3 Slot 4	4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	2	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	2	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module"	2	

Table 20 VNX7500 storage processor configuration (page 5 of 6)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
File and Unified with eight BE buses and two to four blades	Two 4-port SAS I/O modules for BE bus connections	Slot 0 Slot 1	4-port 8 Gb Fibre Channel I/O module	2	Slot 2 Slot 3
	One 4-port Fibre Channel I/O module for blade connections	Slot 4	4-port 1 Gb copper iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	2	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	2	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module”	2	

Table 20 VNX7500 storage processor configuration (page 6 of 6)

Configuration	Required I/O modules per SP		Optional I/O module slots per SP		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per SP	Location
File and Unified with eight BE buses and five to eight blades	Two 4-port SAS I/O module for BE bus connections	Slot 0 Slot 1	4-port 8 Gb Fibre Channel I/O module	1	Slot 2
	Two 4-port Fibre Channel I/O modules for blade connections	Slot 3 Slot 4	4-port 1 Gb copper iSCSI I/O module	1	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	1	
			2-port 10 Gb optical/-TwinAx iSCSI I/O module	1	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module"	1	

IMPORTANT

To upgrade a block system to a two-to-four blade File or Unified configuration after installation, slot 4 in each SP must be left open for a Fibre Channel I/O module. Slots 3 and 4 in each SP must be left open for Fibre Channel I/O modules to upgrade a block system to a five-to-eight blade File or Unified configuration.

Data Mover configuration

Table 21 shows the Data Mover I/O module configurations for the VNX7500 File and Unified configurations.

Table 21 VNX7500 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	4	Slot 1 Slot 2 Slot 3 Slot 4
			2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	4	
			2-port 10 Gb optical Ethernet I/O module	4	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	4	

VNX7600

This section describes the hardware for the VNX7600. The VNX7600 is available in Block, File, and Unified configurations. Figure 34 illustrates a VNX7600 Block configuration.

The VNX7600 uses a DPE which includes a Battery on Bus; as a result, no SPS is necessary. The DPE contains the first 25 disks.




Figure 34 VNX7600 Block configuration

Figure 35 shows a VNX7600 File/Unified configuration rear view.




Figure 35 VNX7600 File/Unified configuration, rear view

Storage processor configuration

The VNX7600 SPs can be upgraded with additional I/O modules for additional FC, SAS, FCoE, or iSCSI connections.

For all configurations, there are 2 SAS ports (0, 1) for BE bus connections. The VNX7600 does support the addition of 4 Back End ports using a SAS I/O module.

The VNX7600 SP supports up to 5 I/O modules in each SP.

Table 22 shows the available I/O module configurations for the VNX7600.

Table 22 VNX7600 storage processor configuration

Configuration	Required I/O module		Optional I/O modules		
	I/O module type	Location	I/O module type	Maximum additional modules supported per SP	Location
Block	No requirement		4-port 8 Gb Fibre Channel I/O module	5 (Block) 4 (File/Unified)	Slots 0-3 if File/Unified Slots 0-4 if Block
File/Unified	4-port 8 Gb Fibre Channel I/O module	Slot 4	4-port 10 GbE iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	5 (Block) 4 (File/Unified)	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	2	
			4-port 1 GBase-T iSCSI I/O module	2	
			4-port 4x lane/port 6 Gb SAS I/O module	1	Slot 0

Note: Both SPs must have the same modules in the same slots.

Data Mover configuration

“Data Mover enclosure (DME)” on page 27 provides more information about the Data Mover hardware.

Table 23 shows the Data Mover I/O module requirements and options for the VNX7600 File and Unified configurations.

Note: The VNX7600 DME supports I/O modules in 4 slots, slot numbers 0-3. Slot 4 has a cover over it as it is not supported.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 23 VNX7600 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1GbBase T for IP I/O module	3	Slot 1 Slot 2 Slot 3
			2-port 10 Gb (optical & active twinax) Ethernet I/O module	3	
			2-port 10 Gb Base-T IP I/O module	3	

Figure 17 on page 28 shows the location of the DME I/O module slots.

VNX8000

The VNX8000 utilizes two 2U Li Ion SPSs. Figure 36 shows the VNX8000 Block configuration. Figure 37 on page 75 shows the VNX8000 Unified configuration.




Figure 36 VNX8000 Block configuration




Figure 37 VNX8000 Unified configuration

Storage processor configuration

The SPE in the VNX8000 has 11 slots for I/O modules. All connectivity to storage disks is handled by ports in SAS I/O modules. The VNX8000 supports either 8-bus systems or 16-bus systems for back end connections. A base system will include I/O modules for up to 8 back end loops in slots 5 and 10 of each SP. Additional I/O modules for front-end connections are optional.

Table 24 Back end connections for a base VNX8000 8 loop system

I/O module in slot #	Bus	Ports
A5, B5	0-3	0-3
A10, B10	4-7	0-3

Because the ports and their order are persisted from initialization (and cannot be easily reconfigured in upgrades), this means that a VNX8000 system that is upgraded in the field [additional SAS modules in slots 4 and 6] will result in the following configuration:

Table 25 Back end loops for VNX8000 system upgraded in the field

I/O module in slot #	Bus	Ports
A4, B4	8-11	0-3
A5, B5	0-3	0-3
A6, B6	12-15	0-3
A10, B10	4-7	0-3

Notice that the ports, cables, etc. for the original SAS modules do not change in the upgrade.

A VNX8000 system with 16-bus back end has port connections different from the base model (because the extra SAS I/O modules and their ports are identified and persisted in order). The 16-bus system shipping from the factory is built this way:

Table 26 VNX8000 Back end connections for 16 loop system

I/O module in slot #	Bus	Ports
A4, B4	1-4	0-3
A5, B5	0, 5-7	0-3
A6, B6	8-11	0-3
A10, B10	12-15	0-3

Table 27 VNX8000 storage processor configuration

Configuration	Required I/O module		Optional I/O modules		
	I/O module type	Location	I/O module type	Maximum modules supported per SP ^a	Location
All	4-port 4x lane/port 6 Gb SAS I/O module	A5, B5 A10, B10			
File/Unified	4-port 8 Gb Fibre Channel I/O module	A8, B8			
			4-port 8 Gb Fibre Channel I/O module	9 (Block) 8 (File/Unified)	
			4-port 10 GbE iSCSI I/O module	2	
			2-port 10 Gb optical/-TwinAx FCoE I/O module	9 (Block) 8 (File/Unified)	
			2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	4	
			4-port 1 GBase-T iSCSI I/O module	2	
			4-port 4x lane/port 6 Gb SAS I/O module	4	

- a. Each SP contains a total of 11 slots for I/O modules. See [Figure 7 on page 20](#). Limits on particular FE modules are subject to the overall system limit. Thus if you have 2 slots with SAS modules for connecting to the array's disks, there are 9 slots left. If you have 4 SAS modules in each SP, then there are 7 other slots total.

The VNX8000 base system provides 1 SAS I/O module in slot 5 and 1 SAS I/O module in slot 10 of each SP. VNX8000 Block configurations can then add I/O modules for front-end connections. In order to balance the workload across the CPU sockets, the following sequence is recommended for front-end connections:

Table 28 Front End I/O Module Sequence

Slot	Priority
8	1
2	2
0	3
9	4
3	5
7	6
1	7
4	8
6	9

Data Mover configuration

[Table 29](#) shows the Data Mover I/O module requirements and options for the VNX8000 File and Unified configurations.

Note: The VNX8000 DME supports I/O modules in all 5 slots of the Data Mover Enclosure.

VNX8000 File/Unified configurations require at least one FC I/O module in each SP for connection to data movers. The first FC I/O module for connection to the data movers is placed in slot 8; the second FC I/O module is placed in slot 2 (per the priority in [Table 28](#)). Additional I/O modules can be added for host connections.

Note: Both SPs must have the same modules in the same slots.

IMPORTANT

All DMEs in a failover group must have the same I/O modules in the same slots.

Table 29 VNX8000 Data Mover configuration

Configuration	Required I/O modules per blade		Optional I/O modules per blade		
	I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
File and Unified	One 4-port FC I/O module for storage processor connections	Slot 0	4-port 1GbBase T for IP I/O module	4	
			2-port 10 Gb (optical & active twinax) Ethernet I/O module	4	
			2-port 10 Gb Base-T IP I/O module	4	

VNX Gateway systems

The VG2 and VG8 are VNX1 hardware systems and run File OE 7.0/7.1. The VG10 and VG50 are VNX2 hardware systems and run File OE 8.1.

VG2

The VG2 contains one or two data movers in a Data Mover Enclosure and one or two Control Stations. See [Figure 38 on page 81](#).

The VG2 connects via Fibre Channel or FCoE SAN to:

- ◆ Symmetrix storage systems
- ◆ CLARiiON storage systems

- ◆ VNX storage systems




Figure 38 VG2 rear view

Data Mover configuration

Table 30 shows the Data Mover I/O module configurations for the VG2 File and Unified configurations.

Table 30 VG2 Data Mover configuration

Required I/O modules per blade		Optional additional I/O modules per blade		
I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
One 4-port Fibre Channel I/O module OR One 2-port FCoE I/O module for SAN connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	3	Slot 1 Slot 2 Slot 3
		2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	3	
		2-port 10 Gb optical Ethernet I/O module	3	
		2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	3	

VG8

The VG8 contains two to eight data movers in 1-4 Data Mover Enclosures and one or two Control Stations. See Figure 39 on page 83 for a rear view of the VG8.

The VG8 connects via Fibre Channel or FCoE SAN to:

- ◆ Symmetrix storage systems
- ◆ CLARiiON storage systems
- ◆ VNX storage systems




Figure 39 VG8 rear view

Data Mover configuration

“Data Mover enclosure (DME)” on page 27 provides more information about the Data Mover hardware.

Table 31 shows the Data Mover I/O module configurations for the VG8 File and Unified configurations

Table 31 VG8 Data Mover configuration

Required I/O modules per blade		Optional additional I/O modules per blade		
I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
One 4-port Fibre Channel I/O module OR One 2-port FCoE I/O module for SAN connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	4	Slot 1 Slot 2 Slot 3 Slot 4
		2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	4	
		2-port 10 Gb optical Ethernet I/O module	4	
		2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	4	

Figure 17 on page 28 shows the location of the DME I/O module slots.

VG10

The VG10 contains one or two data movers in a Data Mover Enclosure and one or two Control Stations. See Figure 40 on page 85.

The VG10 connects via Fibre Channel or FCoE SAN to:

- ◆ Symmetrix VMAX storage systems




Figure 40 VG10 rear view

Data Mover configuration

Table 30 shows the Data Mover I/O module configurations for the VG10 File and Unified configurations.

Table 32 VG10 Data Mover configuration

Required I/O modules per blade		Optional additional I/O modules per blade		
I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
One 4-port Fibre Channel I/O module OR One 2-port FCoE I/O module for SAN connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	3	Slot 1 Slot 2 Slot 3
		2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	3	
		2-port 10 Gb optical Ethernet I/O module	3	
		2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	3	

VG50

The VG50 contains two to eight data movers in 1-4 Data Mover Enclosures and one or two Control Stations. See Figure 41 on page 87 for a rear view of the VG50.

The VG50 is also used for “internal” gateway models--systems that are physically built in a VMAX cabinet using direct connection to the array. See “Symmetrix systems support” on page 88.

The VG50 connects via Fibre Channel or FCoE SAN to:

- ◆ Symmetrix VMAX storage systems




Figure 41 VG50 rear view

Data Mover configuration

“Data Mover enclosure (DME)” on page 27 provides more information about the Data Mover hardware.

Table 33 shows the Data Mover I/O module configurations for the VG50 File and Unified configurations

Table 33 VG50 Data Mover configuration

Required I/O modules per blade		Optional additional I/O modules per blade		
I/O module type and quantity	Location	I/O module type	Maximum additional modules supported per blade	Location
One 4-port Fibre Channel I/O module OR One 2-port FCoE I/O module for SAN connections	Slot 0	4-port 1 Gb copper Ethernet I/O module	4	Slot 1 Slot 2 Slot 3 Slot 4
		2-port 1 Gb copper Ethernet + 2-port 1 Gb optical Ethernet I/O module	4	
		2-port 10 Gb optical Ethernet I/O module	4	
		2-port 10 Gb RJ-45 Base-T iSCSI/IP I/O module	4	

Figure 17 on page 28 shows the location of the DME I/O module slots.

Symmetrix systems support

VNX gateway systems can connect to Symmetrix VMAX and Symmetrix Business Class systems. The VG10 and VG50 can connect only to Symmetrix VMAX arrays. Symmetrix system models include:

- ◆ Symmetrix DMX Family

- ◆ Symmetrix VMAX Series
- ◆ Symmetrix VMAXe Series
- ◆ Symmetrix 3000
- ◆ Symmetrix 5000
- ◆ Symmetrix 8000

The Symmetrix Remote Data Facility (SRDF) manages the failover and failback process between Symmetrix and VNX VG2/VG8 gateway systems.

Refer to the EMC Online Support website, the EMC *E-Lab Interoperability Navigator*, and EMC Powerlink for more information about Symmetrix systems and SRDF .

VG50 internal gateway (VMAX 10K File)

A new VX10VG50 model combines a Symmetrix VMAX 10K and a VNX VG50. They are assembled and cabled in the same cabinet using direct connections to the VMAX 10K engines.

VNX Hardware Upgrades

The following upgrades are available for VNX hardware:

- ◆ Add disk
- ◆ Add DAE
- ◆ Add SPS¹
- ◆ Add I/O module to SP²
- ◆ Add blade
- ◆ Add I/O module to blade
- ◆ Add blade enclosure
- ◆ Add Control Station

In addition, there are procedures to upgrade [swap to faster] I/O modules. Most of these upgrades are procedures available to customers.

-
1. VNX5100 and VNX5300 only.
 2. Not supported on VNX5100.

VNX for Block to Unified upgrades

A user with a VNX Block system can add VNX for File services on the VNX platform. This customer-performable upgrade includes installing the File hardware, installing the operating system (VNX OE for File), and configuring file services within the environment.

IMPORTANT

Systems running VNX OE for Block R31 must use a Services-based procedure for this upgrade. Systems running R32 can order an upgrade kit containing the File HW, the File Enabler SW, and documentation. This upgrade kit is available to customers. VNX2 systems running R33 do not yet support the Block to Unified upgrades.

Assumptions

The following assumptions apply when performing the VNX Block-to-Unified upgrade:

- ◆ VNX OE for File/VNX OE for Block compatible revisions are loaded on the private LUNs as a precondition to the upgrade.
- ◆ The upgrade will provide uninterrupted block data I/O during the file services upgrade, as well as uninterrupted management access.
- ◆ Block data services, MirrorView, SnapView and replication via Recover Point will not be interrupted.
- ◆ Dual Control Station configurations are supported as part of the upgrade.
- ◆ The end user has enough physical space in the cabinet to add the File hardware components.

Phases

The VNX Block-to-Unified upgrade is divided into four phases:

- ◆ Readiness Check - use the Upgrade Readiness Check Wizard to verify and prepare the VNX for Block system for the upgrade
- ◆ Rack & cable - rack and cable the new VNX for File hardware according to the instructions provided and power-up the hardware

- ◆ VNX Installation Assistant - monitor the automated VNX for File software installation and set up the Unified system for configuration
- ◆ Configure - configure the system for production

To successfully complete the VNX Block-to-Unified upgrade:

- ◆ The user verifies that the VNX block system is ready to be upgraded to a Unified configuration using the Upgrade Readiness Check Wizard.
- ◆ The user downloads the VNX for File software packages needed for the upgrade. The download can only be done by using the Unisphere Service Manager.
- ◆ In addition to the network-based software download, the user installs a Block-to-Unified Enabler software package, available on a DVD delivered in the upgrade kit.
- ◆ A two Data Mover upgrade takes 45 minutes or less. The time is calculated from when the VNX for File hardware is powered-on and running to when the Unified configuration is ready for provisioning.

VNX Data in Place Conversions

In-family Data in Place (DIP) conversions allow installed VNX for Block, VNX File, and Unified models to be converted to higher VNX models with a larger maximum number of drives, and increased CPU and memory resources. Depending on the VNX model to be converted, SP pairs may be swapped, DPEs replaced with SPEs, I/O module pairs added, DAEs added, and DMs swapped. These conversions also re-use installed SAS DAEs with associated drives and I/O modules.

DIP conversions begin with the Upgrade Readiness Check (URC) Wizard. The URC Wizard checks to see which software and hardware is needed on an existing VNX model to prepare it for a conversion.

IMPORTANT

The VNX Data in Place conversions do not impact the configuration. One can convert a VNX5300 Block to a VNX7500 Block, for example.

Table 1: VNX Data in Place conversions

Source Platform	Target Platforms				
Model	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
VNX5100	N	Y	Y	Y	Y
VNX5300	N	N	Y	Y	Y
VNX5500	N	N	N	Y	Y
VNX5700	N	N	N	N	Y
VNX7500	N	N	N	N	N

Note: DIP conversions are available for VNX1 systems as shown above. VNX2 systems do not currently support DIP conversions. Conversions from VNX1 to VNX2 systems are not supported.

CHAPTER 2

VNX General SW Configuration Rules

This chapter covers general configuration rules for the storage array. Topics that are specific to a particular software release or level, such as limits, are in succeeding chapters. Topics include:

◆ VNX series MirrorView ports	94
◆ Host connection to VNX series storage systems	95
◆ Fibre Channel storage area network (SAN) configuration rules	95
◆ VNX storage-system port speed rules	98
◆ Mixing Fibre Channel configurations	101
◆ Mixing Fibre Channel, FCoE, and iSCSI configurations	105
◆ iSCSI network configuration rules	110
◆ Virtual provisioning for VNX series storage systems	119
◆ FAST tiering for VNX series storage systems	124
◆ Disk power savings	128
◆ Unisphere and domain connections	128
◆ Booting from SAN guidelines	131
◆ Replicator V2 Incremental Attach	138
◆ Failover configuration rules	140
◆ Non-disruptive software installation (NDU) configuration rules	146
◆ Reserved LUN pool	147
◆ MirrorView/A and MirrorView/S configuration rules	147
◆ SAN Copy configuration rules	150
◆ LUN usage with replication applications	154
◆ VNX Snapshots	155
◆ VAAI support for NFS	157
◆ VASA support for VMware	159
◆ VAAI support for block	160

VNX series MirrorView ports

The MirrorView port location in the VNX family is different from the designation for those ports in EMC CLARiiON arrays. In VNX systems using a DPE with a mezzanine card (VNX5100, VNX5300, and VNX5500), the MirrorView port is the lowest numbered port on that mezzanine card (logical port 0, physical port 2). For systems using a DPE or SPE with available ports in I/O modules, the MirrorView ports may vary depending on the type and number of I/O modules in the storage system.

IMPORTANT

The MirrorView port is established when the system is initialized. It cannot be changed later.

There is a possibility for the port designated as the MirrorView port to be used for other functions (such as connection to data movers) unless the system has been planned for MirrorView use and thus has enough I/O modules in the correct slots.

Table 34 lists the MirrorView ports for factory systems.

Table 34 MirrorView ports for factory systems

Storage system	MirrorView Fibre Channel FE ports	MirrorView iSCSI FE ports
	Physical slot and port number	Physical slot and port number
VNX5100, VNX5300, VNX5500	Mezzanine FC port 0 ^a	Port 0 on the first iSCSI I/O module. Note: The VNX5100 does not support iSCSI.
VNX5200, VNX5400, VNX5600, VNX5700, VNX5800, VNX7500, VNX7600, and VNX8000	MirrorView ports are established at the time of configuration and are the lowest-numbered available ports.	MirrorView ports are established at the time of configuration and are the lowest-numbered available ports.

a. This refers to logical port 0, not physical port 0. The logical port 0 will be identified on the mezzanine card as port 2.

Use Navisphere CLI or Unisphere to determine the ports available for MirrorView.

Host connection to VNX series storage systems

The same server can connect to VNX-series storage systems, *only if*:

- ◆ The server is running the Unisphere Host Agent and/or Unisphere Server Utility for all storage systems.
- ◆ The VNX5100 series storage systems are running Unisphere Manager.
- ◆ The master of the domain with these storage systems is *one* of the following:
 - A VNX5100 storage system is running VNX OE for Block 05.31 or later
 - A VNX5200, VNX5300, VNX5400, VNX5500, VNX5600, VNX5700, VNX5800, VNX7500, VNX7600, or VNX8000 series storage system
 - A Unisphere management station for all storage systems.

Fibre Channel storage area network (SAN) configuration rules

Host connections to the storage processors in VNX systems are by default Fibre Channel connections. Customers can opt to order systems with FCoE-based systems. See [“Disk processor enclosure \(DPE\) and storage processor enclosure \(SPE\)” on page 15](#) and [“VNX configurations by model” on page 30](#).

Fibre Channel host adapter terminology

This document uses the term HBA to refer generically to both Fibre Channel host bus adapters and converged network adapters. Converged network adapters provide FCoE connectivity.

Fibre Channel switch terminology

DWDM - (Dense Wavelength Division Multiplexing). WDM technology uses multiple lasers and transmits several wavelengths of light simultaneously over a single optical fiber. Dense WDM adds significantly more channels.

Fabric - One or more switches connected by E_Ports. E_Ports are switch ports that are used only for connecting switches together. FCoE switches support cascading only through FC E_Ports. You cannot cascade FCoE switches using any Ethernet ports because the switch vendors do not support VE_Ports. The exception to this

rule is a configuration containing the Nexus 4001L switch module used in an IBM Blade Server chassis. In this configuration, you must connect the 4001L switch module directly to the FCoE switch.

ISL - (Inter Switch Link). A link that connects two E_Ports on two different switches.

Path - A path is a connection between an initiator (such as an HBA or CNA port) and a target (such as an SP port in a storage system). Each HBA or CNA port that is zoned to a port on an SP is one path to that SP and the storage system containing that SP. Depending on the type of storage system and the connections between its SPs and the switch fabric, an HBA or CNA port can be zoned through different switch ports to the same SP port or to different SP ports, resulting in multiple paths between the HBA or CNA port and an SP and/or the storage system. Note that the failover software running on the server may limit the number of paths supported from the server to a single storage-system SP and from a server to the storage system. Refer to the [“Path Rules for servers without failover software” on page 114](#).

Single-initiator zoning - Each HBA or CNA port has a separate zone that contains it and the SP ports with which it communicates. EMC recommends single-initiator zoning as a best practice.

Number of servers and storage systems

The number of servers and storage systems are dependent on the available switch ports. However, each server must follow the [“Fibre Channel fan-in rule” on page 97](#) and each storage system must follow the [“Fibre Channel fan-out rule” on page 97](#), using WWPN switch zoning.

Supported Fibre Channel or FCoE switches

For more information, see [“VNX and Fibre Channel Switches” on page 217](#) and the *E-Lab Interoperability Navigator*.

Fibre Channel connectivity

Fibre Channel Switched Fabric (FC-SW) connection to all server types, subject to the restrictions in the VNX sections of the *E-Lab Interoperability Navigator*; Arbitrated Loop Emulation mode (QuickLoop) for connection to FC-AL HP-UX servers. Note that some Fibre Channel switches, such as the Brocade DS4100 switches, do not support QuickLoop.

FCoE connectivity

FCoE connectivity is currently supported for Linux, VMware ESX, AIX, Solaris, and Windows server types, subject to the restrictions in the VNX sections of the *E-Lab Interoperability Navigator*.

Fibre Channel fan-in rule

A server can be zoned to a maximum of four storage systems.

Fibre Channel fan-out rule

The software installed on a VNX series determines the number of servers that can be connected to the storage system. The maximum number of connections between servers and a storage system is limited by the number of initiator records supported per storage-system SP. An initiator is an HBA or CNA port in a server that can access a storage system.

Note that some HBAs or CNA have multiple ports. Each HBA or CNA port that is zoned to an SP port is one path to that SP and the storage system containing that SP. Each path consumes one initiator record. Depending on the type of storage system and the connections between its SPs and the switches, an HBA or CNA port can be zoned through different switch ports to the same SP port or to different SP ports. This results in multiple paths between the HBA or CNA port and an SP and/or the storage system. Note that the failover software environment running on the server may limit the number of paths supported from the server to a single storage-system SP and from a server to the storage system. Refer to the [“Paths from a server to an SP” on page 113](#).

Number of servers

Note: Highly-available (HA) servers contain an equal number of connections to each SP in the storage system.

Initiator records required for MirrorView or SAN Copy

Each path used in a MirrorView or SAN Copy relationship between two storage systems requires one initiator record in *both* storage systems. The number of initiator records required for Data Movers in VNX storage systems is shown in [Table 35](#).

Table 35 Initiator records for Data Movers

VNX system type	Storage-system initiator records required per Data Mover
File, Unified	2
Direct Connect Gateway	2
SAN Gateway	4

VNX storage-system port speed rules

[Table 36](#) lists the port speeds per storage system. The speed of each port on an SP can be set individually using Unisphere or CLI.

Table 36 Storage processor port speeds

Storage system port type	SP port speeds
FC port speeds	2, 4, or 8 Gb/s ^a
SAS port speeds	4 lanes at 6 Gb/s
FCoE port speeds	10 Gb/s

a. The factory default setting is auto (auto negotiate) for a VNX series SP.

Since most ports in VNX storage processors are in I/O modules, see the kinds of I/O modules and their speeds available for each VNX system in [“I/O modules” on page 29](#) and [“VNX configurations by model” on page 30](#).

FCoE support

All VNX Unified systems except the VNX5100 support FCoE I/O modules in the storage processor.

VNX gateway systems use FCoE connectivity to communicate with the VNX for Block or Symmetrix system arrays that support FCoE. This type of communication is known as Initiator Mode.

The FCoE dual port I/O back end I/O module is used in VNX systems. In [Figure 42 on page 99](#), for example, the servers are connected to a VNX for Block or a Symmetrix system by using a VNX gateway, IP switch, and an FCoE switch.




Figure 42 VNX over FCoE gateway configuration

Configuration guidelines

These guidelines apply to the VNX over FCoE configurations:

- ◆ VNX system is supported in Initiator Mode only using one FCoE I/O module.
- ◆ A direct connection between a VNX1 (VG2 or VG8) or VNX2 (VG10 or VG50) gateway to the VNX for Block or Symmetrix or VMAX system without an FCoE switch is not supported.
- ◆ Either 4 Gbps fibre channel, 8 Gbps fibre channel, or 10 Gbps FCoE connections to switches are supported.
- ◆ Optical or TwinAX active cables are required for 10 Gbps switch connections to FCoE switches.
- ◆ A 8 Gbps Fibre Channel I/O module in addition to the 10 Gbps FCoE I/O module may be used for tape backup via switches or directly to a tape backup unit.

- ◆ A maximum of two FCoE I/O modules per array are allowed for the VNX.
- ◆ Use approved 10 Gb Optical or EMC/Brocade TwinAX Active cables and connectors. TwinAx Passive cables are not supported.
- ◆ When FCoE is configured in the VNX gateway, it must be configured in the data mover slot 0 for booting. It may be configured in other slots depending upon the number of switches/arrays in the configuration.
- ◆ When both FCoE and FC are configured in the VNX gateway, it is understood that FC is being used for tape connectivity only.
- ◆ The FCoE system class (for Fibre Channel and FCoE traffic) has a default of 2240 bytes which cannot be changed.
- ◆ For optical connections, use optical cables OM2 (up to 82M) and OM3 (up to 300M).
- ◆ All Optical cables must utilize a 10 Gbs SFP connector scheme with a 10Gb SFP.
- ◆ Use optical cables for Ethernet connections.

FCoE I/O module

The VNX can connect to tape backup devices by using optical cables, either directly or to an FCoE switch.

For an FCoE I/O module in the boot slot, the tape device connects to any two ports of an FC I/O module in a different slot ([Figure 43](#)).

The ports sense the:

- ◆ Speed of the device they connect to and automatically adjust to the appropriate speed (depending on the type of I/O module installed).

- ◆ Topology of the Fibre Channel network they are connected to and automatically configure for arbitrated loop or switched fabric as needed.




Figure 43 FCoE ports

Mixing Fibre Channel configurations

The Fibre Channel (FC) and Fibre Channel over Ethernet (FCoE) front-end ports in a VNX series storage system are unshared. Each unshared SP port has its own independent Fibre Channel host interface with its own unique World-Wide Port Name (WWPN).

Note: An FCoE configuration must be a configuration with one or more single-switch fabrics.

A storage system can connect to n distinct direct attach (FC-AL) configurations or n separate fabrics in the same or different SAN (FC-SW) or FCoE configurations where n equals the number of FC or FCoE front-end ports per system, as listed in [Table 37](#).

Table 37 FC/FCoE front end ports

Storage-system model	Number of FC front-end ports (n)	Number of FCoE front-end ports (n)
VNX5200, VNX5300, VNX5400, VNX5500, VNX5600, VNX5700, VNX5800, VNX7500, VNX7600, VNX8000	Depends on the number and type of optional UltraFlex I/O modules in the storage system, as described in “ VNX configurations by model ” on page 34	

Multiple Fibre Channel configurations for a single server

Depending on the number of host bus adapter (HBA) ports in a server, any server can support multiple configurations, as shown in [Figure 44 on page 103](#). For example, a Solaris server with four HBA ports could be in the following configurations:

- ◆ A direct attached storage configuration with two HBA ports connected to a VNX storage system.
- ◆ A dual-fabric SAN configuration with one HBA port connected to one fabric and one HBA port connected to the other fabric.

Multiple Fibre Channel configurations for a single storage system

Each SP in a storage system can be in multiple instances of the same or different configurations. Therefore a storage system can support simultaneous direct attach (FC only) and SAN configurations, as shown in [Figure 44 on page 103](#).

Example of a server and a storage system in both Fibre Channel SAN and direct attach configurations




Figure 44 Server and storage system in both Fibre Channel SAN and direct attach configurations

Multiple FCoE switch configurations for a single server

Depending on the number of CNA ports in a server, any server can support multiple switch configurations, as shown in [Figure 45 on page 104](#). For example, a Linux server with four CNA ports could be in two dual-fabric SAN configurations, each with one CNA port connected to one fabric, and one CNA port connected to the other fabric.

Multiple FCoE switch configurations for a single storage system

Each SP in a storage system can be in multiple instances of the same or different switch configurations. Therefore a storage system can support simultaneous switch configurations, as shown in [Figure 45 on page 104](#).

Example of a server and a storage system in multiple FCoE switch configurations




Figure 45 Server and storage system in multiple FCoE switch configurations

Mixing Fibre Channel, FCoE, and iSCSI configurations

Fibre Channel, FCoE, and iSCSI storage-system connections for a single server

A single server can connect to the same VNX series through both its FC data ports and its FCoE data ports. As a general rule, a single server cannot connect to a VNX series storage system through *both* its Fibre Channel data ports or FCoE data ports and its iSCSI data ports. The exception to this rule is a server with virtual machines or virtual systems that run instances of the Unisphere or Navisphere Host Agent that are different from the instance running on the kernel system. In this situation, the initiators registered with the storage system by the Navisphere Host Agent for the kernel system and for the virtual machines or systems appear to be from different servers and can therefore be connected to different storage groups. A single server can connect via Fibre Channel HBAs to one VNX system through its Fibre Channel data ports or via FCoE CNAs to one VNX system through its FCoE data ports and via NICs or iSCSI HBAs to another VNX system through its iSCSI data ports.

A server can connect to VNX storage systems with iSCSI data ports and other storage systems with Fibre Channel data ports over a *common network configuration only if* the server:

- Connects to the Fibre Channel storage systems using supported iSCSI bridge devices as shown in [Figure 46 on page 106](#).
- Has a common network configuration, failover software, and driver support for both types of storage systems.

Note: A server cannot be connected to the same storage system through both NICs and iSCSI HBAs.

Example of a server connected to both an iSCSI and a Fibre Channel storage system over a network




Figure 46 Server connected to both an iSCSI and Fibre Channel storage system over a network

Multiple iSCSI configurations for a single server


Depending on the number of NICs or HBAs in a server, any server can support multiple configurations, as shown in [Figure 47 on page 108](#). For example, a Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Solaris, or Linux server with four NICs could be in the following configurations:

- One direct attached storage configuration with two NICs connected to a VNX series storage system.
- One dual-subnet configuration with one NIC connected to one subnet to which one SP in a VNX series storage system is connected, and the other NIC connected to the other subnet to which the other SP in the VNX series is connected.

Multiple iSCSI configurations for a single storage system

Since each iSCSI I/O module in a VNX series SP has at least two iSCSI ports, it can be in multiple instances of the same or different configurations. Therefore, a VNX series storage system can support simultaneous direct attach and network configurations, as shown in [Figure 47 on page 108](#).

Example of a server and a storage system in both iSCSI network and direct attach configurations



Note: The VNX series array ports shown have iSCSI modules. See [“VNX configurations by model” on page 34](#).

Figure 47 Server and storage system in both iSCSI network and direct attach configurations

Multiple iSCSI and Fibre Channel configurations for a single storage system

Since each SP in a VNX series array may have multiple iSCSI and Fibre Channel ports, it can be in multiple instances of the same or different configurations. The storage system can be in iSCSI configurations with one or more servers and in FC configurations with one or more *other* servers. In other words, a VNX series storage system can support simultaneous iSCSI and Fibre Channel configurations with different servers, as shown in [Figure 48 on page 109](#).




Figure 48 VNX storage system with multiple connections

iSCSI network configuration rules

Network connectivity

Dedicated storage networks are highly recommended, but not required. If not using dedicated storage networks, it is recommended that iSCSI traffic be either on separate physical local area network (LAN) segments or on virtual LANs (VLANs) separate from those used for general LAN traffic. We recommend that you do not configure gateways for iSCSI ports, and keep the hosts on the same subnet as the storage-system iSCSI ports.

VLAN tagging is supported for iSCSI ports on VNX storage systems. VLAN tagging supports eight VLAN tags per iSCSI port. Additionally, Ethernet switch-based VLANs are supported and can be useful in isolating iSCSI traffic from general LAN traffic. Servers can connect to a storage system using a Layer 2 (switched) or Layer 3 (routed) network.

VNX series VLAN virtual ports

VLAN virtual ports are supported on VNX storage systems. A VLAN is a group of hosts or devices that lets the devices communicate as if they were attached to a single broadcast domain, regardless of their physical location. A VLAN does not require that you locate the physical devices in close proximity to each other; the devices are configured using the storage-system management software, making them more flexible than physical LANs. A virtual port is a logical representation of a network port. You can assign network port properties to each virtual port, just as you can to each physical port. Using Unisphere, you can create multiple virtual ports for each physical port, so you can create independent logical networks within a physical network, allowing iSCSI data ports to participate in multiple VLANs at the same time. To do this, you must:

- ◆ Configure a trunk port on the participating switch per IEEE 802.1q standards. In this configuration, the switch port passes along the network traffic with some set of VLAN tags and drops all other traffic.
- ◆ Enable VLAN tagging for each virtual port.
- ◆ Assign a VLAN ID to the virtual ports that you want to have IDs. If you assign VLAN IDs, each ID must be unique on each physical port.

Supported network switches and routers

Industry standard network switches and routers with 1-Gigabit copper connectivity or 10-Gigabit optical connectivity to the storage system are supported. Because the 1 Gb/s Ethernet, 10 Gb/s Ethernet iSCSI, and 10 GBase-T iSCSI/IP connection topologies are not interoperable, 1 Gb/s Ethernet, 10 Gb/s Ethernet, and 10 GBase-T iSCSI/IP storage-system iSCSI ports cannot operate on the same physical network. A 1 Gb/s Ethernet network, a 10 Gb/s Ethernet network, and a 10 GBase-T iSCSI/IP network can be connected using a network switch that support both topologies.

If you want to use jumbo frames (that is, a maximum transmission unit of greater than 1500 bytes), all switches and/or routers in the paths to the storage system must support and be set up for the desired maximum transmission unit.

Number of servers and iSCSI storage systems

Servers available on the network can be connected to any iSCSI storage systems available on the network, provided each server follows the following fan-in rule and each storage system follows the following fan-out rule.

iSCSI fan-in rule

A server can be networked to a maximum of four storage systems. A server can have a maximum of four NICs or iSCSI HBAs connected to one storage system. A server *cannot* use both NICs and iSCSI HBAs to connect to SP iSCSI data ports; it must use either all NICs or all iSCSI HBAs. Each initiator can be connected to each SP iSCSI data port in a storage system; however, only one connection per initiator to each SP iSCSI data port is allowed.

HP-UX and Solaris servers supported only NIC connections to VNX for File and VNX for Block storage-system iSCSI data ports. For the latest supported configurations, refer to the *E-Lab Interoperability Navigator* or the Host Connectivity guides on the EMC Online Support website.

A server can connect to VNX for File, VNX for Block, and Symmetrix iSCSI storage systems over a common network configuration, if the server has common failover software and driver support for all storage-system types.

iSCSI fan-out rule

The software installed on a VNX storage system determines the number of iSCSI servers that you can connect to the storage system. The maximum number of connections between servers and a storage system is limited by the number of initiator records supported per storage-system SP.

An initiator is any NIC or HBA port with a *unique* iSCSI name or ID in a server that has access to an SP. The AIX, HP-UX, Solaris, VMware ESX iSCSI initiator software and the Microsoft iSCSI Software Initiator assign the *same* iSCSI name to *all* NIC initiators in the server so they are seen as one single initiator.

Note that some NICs or iSCSI HBAs may have multiple ports. Each NIC or iSCSI HBA port with a connection to an SP port is one path to that SP and the storage system containing that SP. Each path consumes one initiator record. Depending on the type of storage system and the connections between its SPs and the network, a NIC or iSCSI HBA port can be connected through different router or switch ports to the same SP port or to different SP ports. This results in multiple paths between the NIC and the iSCSI HBA port and an SP and/or the storage system. Note that the failover software environment on the server may limit the number of paths supported from the server to a single storage-system SP and from a server to the storage system. Refer to the [“Paths from a server to an SP” on page 113](#).

Number of servers

See [“iSCSI network configuration rules” on page 164](#) for R31-32, and [“iSCSI network configuration rules” on page 191](#) for R33 and later.

iSCSI connection speed

By default, VNX storage systems auto-negotiate the connection speed of their 1 Gb/s Ethernet iSCSI ports. You can force a fixed connection speed of 10, 100, or 1000 Mb/s using either the Unisphere or Navisphere CLI. Ethernet iSCSI ports at a speed of 10 Gb/s in a VNX storage system operate at a fixed rate of 10 Gb/s. Because the 1 Gb/s Ethernet, 10 Gb/s Ethernet iSCSI, and 10 GBase-T iSCSI/IP connection topologies are not interoperable, the 1 Gb/s Ethernet, 10 Gb/s Ethernet iSCSI, and 10 GBase-T iSCSI/IP ports cannot operate on the same physical network. A 1 Gb/s Ethernet network, a 10 Gb/s Ethernet network, and a 10 GBase-T iSCSI/IP network can be connected using a network switch that support both topologies.

iSCSI direct attached storage (DAS) configuration rules

Note: A server *cannot* use both NICs and iSCSI HBAs to connect to iSCSI storage systems; it must use either all NICs or all iSCSI HBAs.

iSCSI connection speed

By default, VNX storage systems auto-negotiate the connection speed of their 1 Gb/s Ethernet iSCSI ports. You can force a fixed connection speed of 10, 100, or 1000 Gb/s using either the Unisphere or Navisphere CLI.

Paths from a server to an SP

Access from a server to a storage processor (SP) in a storage system can be either single path, multipath, or alternate path. Cabling HBA ports, CNA ports, NIC ports, SP ports, and optional switches and routers in specific ways, *enables* paths from a server to an SP. Fibre Channel switch zoning *allows* the paths and *enforces* the paths rules.

Note: Some HBAs and CNAs have two ports. Each port that is zoned to a port on an SP is one path to that SP and the storage system containing that SP. Depending on the type of storage system, the connections between its SPs, and the switch fabric, a CNA's port can be zoned through different switch ports to the same SP or to different SPs. This results in multiple paths between the CNA's port and an SP and/or the storage system. Note that the failover software running on the server may limit the number of paths supported from the server to a single storage-system SP and from a server to the storage system. Refer to the [“Path Rules for servers without failover software” on page 114](#).

Single path - Failover software with single path capability can direct I/O to any specific LUN through only one path at a time. The software can be configured with a maximum of one connection to each SP (or to a single storage system, if using asymmetric active active (ALUA) failovermode). If the active path fails, the failover software redirects I/O through the surviving path to the other SP.

Alternate path - Failover software with alternate path capability can direct I/O to any specific LUN through only one path at a time. Multiple paths can be configured to each SP, but no load-balancing is available. If the active path fails, the failover software redirects I/O through a surviving path, which can be to the same SP or the other SP, depending on where the failure occurred.

Multipath - Failover software with multipath capability can direct I/O to any specific LUN through one or more paths at the same time. If multiple paths are configured to an SP, load balancing is possible (for example, PowerPath distributes I/O using one of several user-selectable algorithms). If an active path fails, the failover software redirects I/O through the surviving paths, which may be to the same SP or the other SP, depending on where the failure occurred. VNX series storage systems support ALUA failover. ALUA is a SCSI Standard which allows the array to describe to the multipath software, on a per LUN basis, which paths will have Optimized performance and which paths are currently Non-Optimized. VNX1 systems would only report, for a single LUN, Optimized paths on one SP. VNX2 systems, in some cases, presents Optimal performance out of all paths (Symmetric Active/Active), and reports that to the multipath software. This may cause the multipath software to use all paths, depending on the version of multipath software.

For information on which failover software for which operating systems support ALUA failover, see Knowledgebase emc99467.

Path Rules for servers without failover software

If a server does not have failover software, only a single path from the server to each SP (or to a single storage system, if using ALUA failovermode) is supported. For more information, refer to [“Supported configurations” on page 140](#).

Note: Without failover software, no SP failover or non-disruptive software installation (NDU) is provided. As a result, while an SP reboots, the server cannot access the LUNs owned by that SP.

Path rules for servers with failover software

Table 38 and Table 39 on page 117 list the number of paths supported from a server with failover software. For information on the storage systems supported by each type of failover software, see “Failover configuration rules” on page 140

Table 38 Number of paths supported from a server (page 1 of 3)

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
AIX	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	Native multipath failover ^c	Multipath	8	2
	DMP ^d	Multipath or Alternate Path ^e	8	16
HP-UX	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	Native multipath failover	Multipath	8	2
	PVLinks	Alternate Path	4	8
	DMP with VNX for Block driver	Multipath or Alternate Path ^e	8	16
Linux	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	Multipath failover (MPIO)	Multipath	4	8
	DMP with VNX for Block driver	Multipath or Alternate Path ^e	8	16
Net Ware	PowerPath	Multipath	N/A	16
	PowerPath SE ^b	Single Path	N/A	2
Open VMS	Native	Multipath	8	16

Table 38 Number of paths supported from a server (page 2 of 3)

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
Solaris	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	DMP with VNX for Block driver	Multipath or Alternate Path ^e	8	16
	Sun StorEdge Traffic Manager (SSTM)	Multipath	8	16
VMware ESX 3, ESXi 3, ESX Server 2.5,	Native	Alternate Path	8	16
Vmware ESX 4, ESXi 4	Native	Multipath	8	16
	PowerPath VE	Multipath	8	16

Table 38 Number of paths supported from a server (page 3 of 3)

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
Windows Server 2008, Windows Server 2008 R2	PowerPath	Multipath	8	16
	Multipath failover (MPIO)	Multipath	8	16

- a. Failover software that supports alternate path or multipath access to an SP also supports single-path access to an SP.
- b. For a description of PowerPath SE, see [“Failover configuration rules” on page 140](#).
- c. All multipath software is not supported for iSCSI configurations, and MirrorView/A, MirrorView/S, or SAN Copy configurations.
- d. See the VERITAS Volume Manager section of the E-Lab Interoperability Navigator for supported configurations.
- e. Multipath for DMP 4.1 or later; alternate path for DMP lower than 4.1.

Table 39 Number of paths supported from a server running Windows Server 2003 and Windows 2000

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
Windows Server 2003	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	DMP with CLARiiON driver	Multipath or Alternate Path ^c	8	16
Windows 2000	PowerPath	Multipath	8	16
	PowerPath SE	Single Path	1	2
	DMP with CLARiiON driver	Multipath or Alternate Path ^c	8	16

- a. Failover software that supports alternate path or multipath access to an SP also supports single-path access to an SP.
- b. For a description of PowerPath SE, see [“Failover configuration rules” on page 140](#).
- c. Multipath for DMP 4.1 or later; alternate path for DMP lower than 4.1.

Sample VNX series configurations

Some examples of valid server connections to a VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500 without MirrorView/A, MirrorView/S, or SAN Copy are:

- Highly available dual-HBA, CNA, or NIC configuration without load-balancing:
 - Each server has two FC HBA, CNA, iSCSI HBA, or NIC ports.
 - Each FC HBA, CNA, iSCSI HBA, or NIC is cabled and zoned or setup to allow it a path to only one port on one SP.
 - Since each server has two FC HBA, CNA, iSCSI HBA, or NIC ports, each server must support two paths to the VNX system.
 - Initiator records used per storage system equals the number of servers x 2 paths.

For the maximum number of servers and initiator records for the VNX systems, see either [“Fibre Channel fan-out rule” on page 97](#) or [“iSCSI fan-out rule” on page 112](#).

- Highly available dual-HBA, CNA, or NIC configuration with maximum load-balancing:
 - Each server has two FC HBA, CNA, iSCSI HBA, or NIC ports.
 - Each FC HBA, CNA, iSCSI HBA, or NIC port is cabled and zoned or setup to allow it a path to four ports on each SP.
 - Since each server has two FC HBA, CNA, iSCSI HBA, or NIC ports, each server must support 16 paths to the VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500.
 - Initiator records used per storage system equals the number of servers x 16 paths.

For the maximum number of servers and initiator records for the VNX storage series, see either [“Fibre Channel fan-out rule” on page 97](#) or [“iSCSI fan-out rule” on page 112](#).

Some examples of valid server connections to a VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500 with MirrorView/A or MirrorView/S and SAN Copy are:

Note: MirrorView/A, MirrorView/S, and SAN Copy are not supported for FCoE configurations.

- Highly available dual-HBA or NIC configuration without load-balancing:
 - Each server has two FC HBA, iSCSI HBA, or NIC ports.
 - Each FC HBA, iSCSI HBA, or NIC is cabled and zoned or set up to allow it a path to only one port on one SP.
 - Since each server has two FC HBA, iSCSI HBA, or NIC ports, each server must support two paths to the VNX series storage system.

- MirrorView / A or MirrorView / S connections from the VNX series system to one other storage system.
- SAN Copy connections from one port on each SP on the VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500 to one port on each SP on another storage system without SAN Copy.
- Initiator records used per storage system equals the number of servers x 2 paths + 2 MirrorView / A or MirrorView / S initiator records + 2 SAN Copy initiator records.

For the maximum number of servers and initiator records for the VNX storage series, see either [“Fibre Channel fan-out rule” on page 97](#) or [“iSCSI fan-out rule” on page 112](#).

- Highly available dual-HBA or NIC configuration with maximum load-balancing:
 - Each server has two FC HBA, iSCSI HBA, or NIC ports.
 - Each HBA port is cabled and zoned to allow it a path to all four ports on each SP.
 - MirrorView / A or MirrorView / S connections from the VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500 to four other storage systems.
 - SAN Copy connections from two ports on each SP on the VNX series system to two ports on each SP on another storage system.
 - Since each server has two HBA ports, each server must support 16 paths to the VNX series system.
 - Initiator records used per storage system equals the number of servers x 16 paths + 8 MirrorView / A or MirrorView / S initiator records + 4 SAN Copy initiators.

For the maximum number of servers and initiators for the VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500, see either [“Fibre Channel fan-out rule” on page 97](#) or [“iSCSI fan-out rule” on page 112](#).

Virtual provisioning for VNX series storage systems

Virtual provisioning for VNX series storage systems lets you assign storage capacity to a server by using pools on which you create thick and/or thin LUNs.

Pools

A set of disks with the same Redundant Array of Independent Disks (RAID) type (RAID 6, RAID 5, or RAID 1/0) that shares its user capacity with one or more pool (thick or thin LUNs). RAID 5 is the default RAID type for a pool. For best practices, use RAID 6 for less efficient drives and RAID 5 or RAID 1/0 for more efficient drives. You can intermix

different types of disk – Flash, SAS, and NL-SAS – in a pool. The disks in a pool cannot be vault disks (000-003). We recommend that all the disks of the same type in a pool have the same capacity.

Pool LUN

A set of disks with the same RAID type that look like an individual disk to an operating system on a server connected to the storage system. Two types of pool LUNs are available – LUNs without the thin property, called *thick LUNs*, and LUNs with the thin property, called *thin LUNs*. LUNs on pools are often called *pool LUNs* to distinguish them from LUNs on RAID groups (RAID group LUNs).

IMPORTANT

With VNX2 arrays, the default LUN selection type, when creating Pool LUNs, is Thin.

You can perform the following operations on a pool LUN:

- ◆ Expand the LUN's capacity without using metaLUNs
- ◆ Shrink the LUN's capacity
- ◆ Compress the data on the LUN if the storage system has the Compression enabler installed
- ◆ Convert LUNs from thick to thin or thin to thick by using compression rather than migration
- ◆ Auto-tier the LUN if the storage system has the FAST enabler installed

Thick LUNs

The capacity of a thick LUN, like the capacity of a RAID group LUN, is distributed equally across the disks in the pool on which you create the LUN. The amount of physical space allocated to a thick LUN, like the capacity of a RAID group LUN, is the same as the user capacity seen by the server's operating system. A thick LUN uses slightly more capacity than the amount of user data written to it due to the metadata required to reference the data. Unlike a thin LUN, a thick LUN can never run out of space.

Thin LUNs

Thin LUNs are supported for storage systems with the Thin Provisioning enabler installed. A thin LUN competes with other LUNs in the pool for the pool's available storage. The capacity of the thin LUN that is visible to the server is independent of the available physical storage in the pool. To a server, a thin LUN behaves very much like a RAID group LUN or a thick LUN. Unlike a RAID group LUN or a thick LUN, a thin LUN can run out of disk space if the pool to which it belongs runs out of disk space. Such an event is an unrecoverable write error and data from the last write operation will not be available. By default, the storage system issues a warning alert when 70% of the pool's space has been consumed. A critical alert is issued when 85% of the space has been consumed. You can customize these thresholds that determine when the alerts are issued. As thin LUNs continue consuming the pool's space, both alerts continue to report the actual percentage of consumed space. A thin LUN uses slightly more capacity than the amount of user data written to it due to the metadata required to reference the data.

Table 40 lists virtual and traditional provisioning for VNX OE for block features.

Table 40 Virtual provisioning and traditional provisioning (page 1 of 4)

Feature	Virtual provisioning		Traditional provisioning
	Thick LUN	Thin LUN	RAID group LUN
Pool RAID types	RAID 6, RAID 5, or RAID 1/0	RAID 6, RAID 5, or RAID 1/0	RAID 6, RAID 5, RAID 3, RAID 1/0, or RAID 1 RAID groups, individual disk, or hot spare.
LUN expansion	Fully supported	Fully supported	Only supported using metaLUNs
LUN shrinking	Fully supported for Windows Server 2008 and Windows Server 2008 R2 hosts.	Fully supported for Windows Server 2008 and Windows Server 2008 R2 hosts.	Fully supported for Windows Server 2008 and Windows Server 2008 R2 hosts connected to a storage system running R30 or later for pool LUNs (thin and thick) and R29 for RAID Group LUNs.
LUN compression	Fully supported by enabling block compression without performing a migration if the Compression enabler is installed.	Fully supported by enabling block compression without performing a migration if the Compression enabler is installed.	Fully supported for RAID group LUNs by enabling block compression without performing a migration. Supported if the Compression enabler is installed.

Table 40 Virtual provisioning and traditional provisioning (page 2 of 4)

Feature	Virtual provisioning		Traditional provisioning
	Thick LUN	Thin LUN	RAID group LUN
LUN migration	Fully supported	Fully supported	Fully supported
LUN conversion	Fully supported	Fully supported	Use instead of LUN compression, where a Thick pool LUN is converted to a Thin pool LUN or a Thin pool LUN is converted to a Thick pool LUN
Disk usages	Any type of disk, including Flash (SSD) disks, can be in the pool with the thick LUN. The disks in the pool with the thick LUN cannot be the vault disks 000–003.	Any type of disk, including Flash (SSD) disks, can be in the pool with the thin LUN. The disks in the pool with the thin LUN cannot be the vault disks 000–003.	All disks in the RAID group with the LUN must be of the same type.
Space efficiency	When you create a thick LUN, the LUN is assigned physical space on the pool equal to the LUN's size. This space is always available to the LUN even if it does not actually use the space.	When you create a thin LUN, a minimum of 2 GB of space on the pool is reserved for the thin LUN. Space is assigned to the thin LUN on an as-needed basis. Since the thin LUNs compete for the pool's space, a pool can run out of space for its thin LUNs. Such an event is an unrecoverable write error and data from the last write operation will not be available. Some applications, such as Veritas Storage Foundation, may return space no longer needed to the pool by an appropriate signaling to the storage system.	When you create a LUN, the LUN is assigned physical space on the RAID group equal to the LUN's size. This space is always available to the LUN even if it does not actually use the space.

Table 40 Virtual provisioning and traditional provisioning (page 3 of 4)

Feature	Virtual provisioning		Traditional provisioning
	Thick LUN	Thin LUN	RAID group LUN
Performance	Thick LUN performance is comparable to the performance of a RAID group LUN and is typically faster than the performance of a thin LUN.	Thin LUN performance is typically slower than thick LUN performance.	RAID group LUN performance is comparable to the performance of a thick LUN and is typically faster than thin LUN performance.
Manual administration	Pools require less manual administration than RAID groups.	Pools require less manual administration than RAID groups.	RAID groups require more manual administration than pools.
Use with SnapView	Fully supported for thick LUNs. A thick LUN cannot be a clone private LUN.	A thin LUN can be a snapshot source LUN, a clone LUN, a clone source LUN, but not a clone private LUN, and not in the reserved LUN pool. A thin LUN cannot be a clone private LUN.	Fully supported for RAID group LUNs.

Table 40 Virtual provisioning and traditional provisioning (page 4 of 4)

Feature	Virtual provisioning		Traditional provisioning
	Thick LUN	Thin LUN	RAID group LUN
Use with MirrorView/A or MirrorView/S	Fully supported for thick LUNs. A thick LUN cannot be used in the write intent log.	Mirroring with thin LUNs as primary or secondary images is supported only between storage systems running VNX OE for Block 05.31 or later. For mirroring between storage systems running VNX OE for Block 05.32, the primary image, secondary image, or both images can be thin LUNs. A thin LUN cannot be used in the write intent log.	Fully supported for RAID group LUNs.
Use with SAN Copy	Fully supported for thick LUNs in all configurations.	Thin LUNs are supported only for SAN Copy sessions in the following configurations: Within a storage system running VNX OE for Block 05.31 or later Between systems running VNX OE for Block 05.31 or later The source LUN must be on the storage system that owns the SAN Copy session.	Fully supported for RAID group LUNs in all configurations.

FAST tiering for VNX series storage systems

Storage tiering lets you assign different categories of data to different types of storage to reduce total storage costs. You can base data categories on levels of protection needed, performance requirements, frequency of use, costs, and other considerations. The purpose of tiered storage is to retain the most frequently accessed and most important data on fast, high performance (most expensive) disks. The less frequently accessed and less important data can be moved to low performance (less expensive) disks.

Within a storage pool that is not a RAID group, storage from similarly performing disks are grouped together to form a *tier* of storage. For example, if you have Flash (SSD) disks, SAS disks, and NL-SAS disks in the pool, the Flash disks form one tier, the SAS disks form a tier, and the NL-SAS disks form another tier. Based on your input or internally computed usage statistics, portions of LUNs (slices) can be moved between tiers to maintain a service level close to the highest performing storage tier in the pool, even when some portion of the pool consists of lower performing (less expensive) disks. The tiers from highest to lowest are Flash, SAS, and NL-SAS.

Fully Automated Storage Tiering (FAST) is not supported for RAID groups because all the disks in a RAID group, unlike in a pool, must be of the same type (all Flash, all FC, or all SATA). The lowest performing disks in a RAID group determine a RAID group's overall performance.

Two types of tiered storage are available:

- Initial tier placement
- Auto-tier placement

Initial tier placement

The initial tier policies available for storage systems that do *not* require the FAST enabler are listed in [Table 41](#).

Table 41 Initial tier settings for LUNs in a pool (FAST enabler not installed)

Initial tier placement policy	Description
Optimize for Pool Performance (default)	No tier setting specified.
Highest Tier Available	Sets the preferred tier for data relocation to the highest performing disk drives with available space.
Lowest Tier Available	Sets the preferred tier for data relocation to the most cost effective disk drives with available space.

Initial tier placement requires manual specification of the storage tier on which to initially place the LUN's data. After initial tier replacement, either manual compression of the LUN to relocate the data to a different tier or installation of the FAST enabler, which once installed, performs the compression automatically.

Auto-tier (FAST) policies

FAST policies are available for storage systems with the FAST enabler installed. The FAST feature automatically compresses the data between storage tiers to provide the lowest total cost of ownership. Pools are configured with different types of disks (Flash/SSD, SAS, and NL-SAS). The storage-system software continually tracks the usage of the data stored on the LUNs in the pools. Using these LUN statistics, the FAST feature relocates data blocks (slices) of each LUN to the storage tier that is best suited for the data, based on the policies listed in [Table 42](#).

Table 42 FAST Policies for LUNs in a pool (FAST enabler installed)

FAST policy	Description
Auto-tier	Sets the initial data placement to Optimized Pool and then relocates the LUNs data based on the LUN performance statistics such that data is relocated among tiers according to I/O activity.
Highest Tier Available	Moves data to the highest tier available.
Start High, then Auto-tier	Moves data to the highest tier available with subsequent tiering based on activity level.
Lowest Tier Available	Moves data to the lowest tier available.
No Data Movement	Moves no data between tiers, and retains the current tier placement.

If you install the FAST enabler on a storage system with an initial tier placement setting specified, the storage-system software bases the FAST policies on the initial tier policies, as shown in [Table 43](#).

Table 43 Interaction between initial tier placement settings and FAST policies

Initial tier placement before FAST enabler installed	Default FAST policy after FAST enabler installed
Optimize for Pool Performance	Auto-Tier
Highest Tier Available	Highest Tier Available
Lowest Tier Available	Lowest Tier Available
n/a	No Data Movement (retains the initial tier placement settings)

Data relocation scheduling

The Unisphere or Navisphere CLI lets you schedule the days of the week, start time, and durations for data relocation for all participating tiered pools in the storage system. The Unisphere or Navisphere CLI also lets you initiate a data relocation at any time manually. To ensure that up-to-date statistics and settings are accounted for properly prior to a manual relocation, FAST analyzes all statistics gathered independently of its regularly scheduled hourly analysis before starting the relocation.

Load balancing

Load balancing redistributes slices across all drives in a tier of a pool to improve performance. In proactive load balancing, the slices are relocated from one RAID group in a tier to another to balance the load within the tier, based on activity level. These relocations occur within the user-defined tiering relocation window.

Load balancing is critical for pool performance, since expanding pools requires users to either expand the pool with the same number of drives of the original pool size (which contradicts efficiency and may not be possible based on the pool and system configuration), or to expand by a small number of drives (which causes poor performance due to a substantially reduced spindle count for new data going to the newly added drives). This is required for competitive feature parity.

Load balancing is useful in expanding given tier(s) in pools while maintaining or improving pool performance, and removing hot spots in a pool.

Resources

Additional information on tiered storage, including information on location of disk types, optimizing the system for performance, etc. is available in a number of other publications including:

- ◆ VNX hardware information guides
- ◆ *VNX Best Practices for Performance Applied Best Practices Guide*
- ◆ VNX-F Tech Note

Disk power savings

Some disks have a power savings (spin down) option, which lets you assign power saving settings to a RAID group comprised of these disks in a storage system. If power saving is enabled for both the storage system and the RAID group, the disks in the RAID group will transition to the low power state after being idle for at least 30 minutes. Power savings is not supported for a RAID group if any LUN in the RAID group is participating in a MirrorView/A, MirrorView/S, SnapView, SAN Copy session, or metaLUNs. Background verification of data (sniffing) does not run when disks are in a low power state. However, every 24 hours, any RAID group that is in a low power state is powered on and background verification starts at that time. In essence, the background verification is done in more of a batch mode when RAID groups are in a low power state rather than its standard routine when they are powered on. For the currently available disks that support power savings, refer to the EMC® *VNX Disk and OE Matrix* on the EMC Online Support website

Unisphere and domain connections

VNX systems use Unisphere to manage the storage system. Unisphere allows a user to manage multiple systems using the concept of a storage domain, including legacy (CLARiiON, Celerra) and other VNX systems.

Unisphere provides integrated storage domains so that administrators can easily log in to multiple systems.

VNX systems began using Unisphere v1.1. This version does not support integrated storage domains. Unisphere v1.1.25 and later do support them. Legacy systems cannot be members of a VNX local domain. They must be configured as a multi-domain.

Table 44 Unisphere feature support

Feature	Unisphere v1.1	Unisphere v1.1.25	Unisphere v1.2 and later
Management of previous generations systems	N	Unisphere v1.1.25 and later can only manage Celerra systems with DART 6.0 and CLARiiON systems running FLARE 23 to FLARE 30.	Unisphere v1.1.25 and later can only manage Celerra systems with DART 6.0 and CLARiiON systems running FLARE 23 to FLARE 30.
Integrated storage domains, multi-domains, and single sign-on	N	Y	Y
Aggregated alerting	N	Y	Y
Dashboard	Y	Y	Y
Navigation and user interface	Y	Y	Y
System dashboard	Y	Y	Y
Tables	Y	Y	Y
Hardware Views	Y	Y	Y
RecoverPoint SE integration	Y	Y	Y
Link and launch products	Y	Y	Y
Support eco-system	Y	Y	Y
Persistent settings	Y	Y	Y
Certificate validation	Y	Y	Y
Updated global user roles	Y	Y	Y
Unisphere Client and Server packages	Y	Y	Y
MirrorView and SAN Copy management from GUI	N	Y	Y

Table 44 Unisphere feature support (continued)

Feature	Unisphere v1.1	Unisphere v1.1.25	Unisphere v1.2 and later
Updated Support Assets	N	N	Y
Updated Storage System Configuration Display (XML)	N	N	Y
Unified Network Services	N	N	New feature in VNX OE for Block R32 and VNX OE for File 7.1. SPs running R32 and earlier are supported. However Control Stations running DART or VNX for File 7.0 and earlier are not supported. VNX gateway systems are not supported.

Support of multi-domain environments and local users

Like Unisphere 1.0, Unisphere 1.1.25 supports multi-domain environments and local users.

A multi-domain environment lets you manage and monitor a group of domains (potentially all the systems in your storage enterprise) using the same instance of Unisphere.

Local user accounts are used to manage features on the local system. Logging into a system using a local user account is recommended when there are a large number of systems in the domain and you want to restrict visibility to a single system. Using a local user account is also useful when logging into storage processors (SPs) to perform special operations (local user for block) or the Control Station to perform special operations only available to the root and nasadmin accounts (local user for file). You create local user accounts with the privileges associated with the new, consolidated roles introduced in Unisphere 1.1.

You can use local user accounts to manage all system features if you create the same user account and password in both the block and file portions of the system.

Permit asynchronous login

The current release of Unisphere and USM supports asynchronous login for Unified and File-only VNX systems. Asynchronous login allows you to successfully log into your system even if either the file or block portion of the system is not available. A system to which you are partially logged into is referred to as operating in degraded mode because the system content associated with the portion of the system that is not available is grayed out. This can occur to a global user login due to network issues or invalid or rejected certificates. This can occur to a local user login if identical local user accounts do not exist on both the file and block portion of the system or if you have identical local user accounts but one portion of the system is not currently available

Retain domain security settings on exit

Unisphere 1.1.25 requires that VNX systems must be initialized before they can be managed by Unisphere. The user is prompted to initialize security if connecting to a system with un-initialized security and the GUI will exit if the user chooses not to create a global user. The current behavior for the Setup page and CLI for Block is unchanged and does not require that security be initialized.

In addition, when you remove a system from a domain, the system is automatically re-initialized and security is reset to the default user account sysadmin. The password is also reset to the default sysadmin unless you change it when prompted. EMC recommends that you change the default password (and record the new password for future reference).

For additional resources on Unisphere and its features, see:

- ◆ the version of Unisphree online help for your system
- ◆ *EMC Unisphere: Unified Storage Management Solution for the New VNX Series VNX5200, VNX5400, VNX5600, VNX5800, VNX7600 & VNX8000 A Detailed Review* [VNX2 white paper]
- ◆ *Domain Management with EMC Unisphere for VNX* [VNX2 white paper]

Booting from SAN guidelines

VNX storage systems that support for booting from SAN are listed in the *E-Lab Interoperability Navigator*. The *E-Lab Interoperability Navigator* contains the most up-to-date configuration information on supported products and software

requirements. Detailed step-by-step installation procedures are available in the HBA or CNA installation guides, which are available on the HBA or CNA vendor websites in the EMC section.

General considerations

This section covers high-level guidelines and answers to commonly asked questions. For operating system specific guidelines, refer to [“Operating-system-specific boot information” on page 135](#).

Maximum boot servers

To ensure high performance of boot LUNs, pools or RAID groups, boot servers must be configured with sufficient physical disk resources to handle the number of servers that may be booting simultaneously. Each physical disk making up a pool or RAID group used for booting from SAN can support multiple servers. The specific guidelines for servers per disk depends on both the type of physical disk and the boot LUN type, as shown in [Table 45 on page 132](#).

Table 45 Maximum boot servers per pool or RAID group consisting of n devices

Disk Type	Boot LUN type		
	RAID group	Pool of thick LUNs	Pool of thin LUNs
FC rotating disk	5n	5n	2n
Flash disk (SSD)	10n	10n	4n

All disks actively servicing I/O are counted. For instance, a 4+1 RAID 5 has a total of five disks actively servicing I/O; a RAID group of five Flash disks, each capable of handling 10 servers (10n in [Table 45 on page 132](#)) supports a total of 50 boot servers. An eight disk RAID 1/0 pool of thick LUNs has a total of eight disks, and since each rotating disk can support five servers (5n in [Table 45 on page 132](#)), the pool supports a total of 40 boot servers. These recommendations are conservative guidelines, and application load is the main factor in overall performance.

Server boot LUNs

EMC recommends that the boot LUN *not* be in a RAID group containing the vault disks (disks 0-3 or 0-4 on loop 0 in enclosure 0). These disks contain the storage-system boot devices. Current VNX series storage systems do *not* support booting from LUNs comprised of ATA drives. Booting is not supported from a RAID group with power savings enabled.

Failover Mode

Using storage-system ALUA failover mode is especially advantageous for booting from SAN because with ALUA the boot LUN is accessible from either SP, which allows booting to occur when a failure affects the primary boot path. ALUA failovermode is the default failovermode on VNX series storage systems.

Pagefiles location considerations

In some operating systems, such as Windows and AIX, the default installation location of the pagefile (also called the swapfile) is on the same LUN as the boot partition (rootvg in AIX). In other operating systems, such as Solaris, the user must set up a configuration file specifying the device to mount on boot and the location of the pagefile. Refer to the HBA configuration guides for details regarding pagefile setup. Some general considerations for deciding whether to locate the pagefiles on internal drives, on storage-system LUNs, or both, are provided below. All of the configurations below are helped by failover and multipath software, such as PowerPath. The issues detailed below are generic and do not assume the additional benefits of such software.

Local pagefile

Performance

Pagefiles have small I/O and high frequency. As such, performance is a key consideration in the location of pagefiles. Locating pagefiles locally may improve the performance, and with the frequency of use, a local location could yield significant improvement in system performance.

No benefit to replication

Locating boot LUNs in a storage system is a benefit when using storage-system-based replication tools to copy these LUNs; that is, for creating SnapView clones for local replicas or MirrorView mirrors for remote replicas. Having boot LUNs in a storage system provides no benefit for replicating Pagefile information because this data is not used to boot the secondary server. Refer to

subsequent notes on location of pagefile information relative to the boot partition and the replication of boot LUNs.

Memory dump file limitations (Windows)

Locating the boot partition on a storage-system LUN and the pagefile on the internal drive is counter to the default installation location of the pagefile for some operating systems. In Windows, this configuration prevents the memory dump file from being created when a memory dump occurs because the system memory is not able to write to the pagefile since it is not in the default location. As a result, only locating the pagefile remotely in the same location as the boot partition allows the creation of dump files. For additional information on this behavior in Windows environments, refer to Microsoft KnowledgeBase Article # 305547 - *Support for Booting from a Storage Area Network (SAN)*.

Remote pagefile

Support for diskless servers

A major reason to locate the pagefiles remotely is to have servers without internal drives, and have the entire operating system environment located on the storage system.

Equivalent robustness

If the boot device is located remotely, locating the pagefiles on the internal servers helps only minimally should a loss of connectivity to the storage system happen because a server error can occur if the boot partition files become inaccessible. Consequently, locating the pagefile internally may assist with performance and make the server less sensitive to path loss, but it does not make the overall environment any more robust.

Memory dump file support (Windows)

As noted previously, for some operating systems, such as Windows, the default installation location of the pagefile is on the same LUN as the boot partition. For these operating systems, if the boot partition is located on a storage-system LUN, the pagefile is installed on the same LUN by default. This enables the creation of the memory dump file because the system memory is written to the same location as the pagefile when a memory dump occurs. Locating the pagefile remotely is thus a requirement for the creation of dump files in these operating systems.

Combination of Local and remote pagefile locations

Considering the benefits of locating the pagefiles locally or remotely, users may want to create multiple pagefiles, and store some locally and some remotely. In this case, the remote pagefile supports the creation of a dump file because it is located on the same LUN as the boot partition, while the local pagefile may increase performance.

Other general notes on pagefiles

Operating-system and/or hardware support

Some server BIOS and/or server vendors simply do not support having local disks installed if you want to boot externally from installed HBAs. Additionally, some servers do not allow the internal controller to be disabled, only allow it to be bypassed if it does not control any disks. Refer to vendor specifications for more details.

Pagefile sizing

The size of the pagefiles is an additional consideration in some operating systems. For example, to enable the saving of crash dump files in Windows, the pagefile on the boot partition must be the same size as physical RAM plus 1 MB to write out the dump file. Typically Windows issues a warning if the pagefile is not set to this minimum. Refer to the HBA configuration guide for specific details regarding pagefile sizes.

Pagefiles and replication

Users should keep in mind that if they want a local pagefile and also want to replicate the boot LUN, the secondary environment must be configured identically. For example, in a Windows environment, the secondary server must be configured to have the same drive letter for the internal drive, so that Windows knows to look for the pagefile in the same location. Any changes in the secondary server configuration may make the server unable to boot from the replicated LUN. Additionally, the user should ensure they have proper licensing for any operating system or storage-system software if using the secondary LUNs.

Operating-system-specific boot information

For details on supported attaches, operating-system-specific guidelines, and boot limitation guidelines, refer to the following:

- *E-Lab Interoperability Navigator* at <http://elabnavigator.emc.com>.
- Host connectivity guides and HBA installation guides on the EMC Online Support website.
- The operating-system-specific notes below.

AIX boot information

EMC does not support booting off clone devices because they are made not available during configuration or system boot.

HP-UX (and OpenVMS) boot information

Refer to the HP-UX connectivity guide.

Linux boot information

The Linux host connectivity guide has a single section for both VNX for Block and Symmetrix systems for booting from SAN.

Solaris boot information

On the HBA vendor websites below, refer to the sections that reference SAN boot or Fibre Channel boot. These websites refer to the Sun website for installation and configuration instructions for the x86/Sparc common (Leadville) driver. In these instructions, refer to the “SAN boot” section. The Solaris host connectivity guide also contains instructions for boot from SAN.

For all Emulex HBAs: <http://www.emulex.com/downloads/emc.html>

For for the QLogic QLA2462 HBA:

http://support.qlogic.com/support/oem_product_detail.asp?p_id=937&oemid=65&oemname=QLA2462

Note: The Emulex website is arranged differently than the QLogic website.

VMware

In the VMware host connectivity guide, refer to the section on booting from a VNX for Block.

Windows

In the Windows host connectivity guide, refer to the section titled “Bootting Windows from external storage section and General Host Behavior and Limitations.”

Replication and boot LUNs

Bootting is supported from a SnapView (clone/snapshot) and remote replication (MirrorView and SAN Copy) replica with the following considerations:

Note: MirrorView and SAN Copy are not supported over FCoE connections.

- Active replicas may create additional loading and response delay on the source LUN. This should be monitored using the procedures noted in the EMC host connectivity guides referenced in the host information sections below. Note that operating system mirroring is not trivial if you consider all the necessary preparation and potential active states in which a host may fail and leave incomplete data on the boot LUN.
- EMC will approve replication of the operating system and data to the secondary LUN, but will *not* be responsible for any re-configuration of hosts, HBAs, operating system, filesystem repair and any additional LUN mappings that applications may require. As with any boot-LUN replication, since the host maintains active operating system data in memory, if the replication is interrupted by primary host failure before it flushes data to LUNs, there is no guarantee that secondary target LUN will be in a consistent or operating-system-usable state. If SnapView is used, the secondary snapshot LUN may not contain complete or usable data if the source LUN is unavailable.
- Users are also responsible to investigate any application or operating-system licensing requirements for utilizing mirrored data copies. Secondary server hardware must be identical to the primary boot server. Users must also adhere to all other boot guidelines and limitations for EMC storage.

Support for boot LUNs with clusters

Booting from SAN is supported for all cluster operating systems listed in the “Clustered Host” sections of the *E-Lab Interoperability Navigator* (subject to any cluster operating-system specific restrictions). For detailed configuration recommendations and requirements, users should consult the operating system cluster guidelines and the *EMC Host Connectivity Guide* for their operating system. Some general notes regarding booting clusters from SAN are provided below, followed by some details specific to each clustered operating system.

Boot LUN isolation

Users should be aware that cluster environments present an additional level of complexity, in that each node of the cluster must uniquely see its boot devices and only its boot devices. Boot devices cannot be shared. At the same time, however, the common purpose of clustered servers is to have shared access to the additional data LUNs.

Single storage-system implementation - Environments with only a single storage system require independent storage groups for each node in the cluster. This

gives each node access to the storage system as a non-clustered server. For instance, in an N-way cluster, users can accomplish this in one of the following ways:

- Create N independent storage groups, as described above, with each appropriate boot LUN in each of the N storage groups. In addition, put the data LUNs in each of the N storage groups, so that they are shared. Unisphere prompts the user to ensure that this is the intended configuration, but in this case, since the configuration is intended, users can disregard the warning.
- Create N+1 storage groups with one storage group containing all the shared cluster LUNs and N Storage Groups each with only the individual boot LUN. This requires manual initiator registration to isolate the boot HBA to only the boot LUN storage group for each server.

Multiple storage-system implementation - Environments with two storage systems have a third option for accommodating the isolation of the boot LUNs. In this scenario, one storage system contains the independent storage groups with the appropriate boot LUNs in each storage group, and the other storage system contains the shared storage group with the shared data LUNs.

Cluster software

HACMP

Supported for booting from SAN, provided that the guidelines given above regarding isolation of boot LUNs from data LUNs are followed.

Sun Cluster 3.x

Currently not supported for boot.

Veritas Cluster Software (VCS)

Currently not supported for boot.

Replicator V2 Incremental Attach

Replicator V2 Incremental Attach provides file system and Virtual Data Mover (VDM) incremental attach for the following two file system replication configurations:

- ◆ Cascading replication, a -> b -> c:
drop b, then establish a -> c incrementally
- ◆ 1 to n replication, a -> b and a -> c:
drop a, then establish c -> b incrementally

The use of Replicator V2 Incremental Attach avoids a time consuming and costly Network Data Management Protocol (NDMP) tape transport or WAN-based full data copy of all replicated file systems when replacing either the source or target VNX with a next generation platform.

Use Cases

Use Case 1: Allow a hardware refresh/replacement of a destination VNX using cascading replication (VNX B and VNX C are at the destination site) to avoid a full copy of all remotely replicated file systems on VNX A at the source site over the WAN to the new target VNX C.

(Source A --> Old Destination B --> Destination Refresh C) ==> (Source A --> Destination C)

Use Case 2: In the event of a Data Replicator site failure, allow the source VNX to quickly reestablish a connection using incremental remote replication with Data Replicator site 2.

(Source A --> DR1 B --> DR2 C) ==> (Source A --> DR2 C)

Use Case 3: In the event of a source site failover and loss of VNX A, allow VNX B and VNX C to reestablish a connection using incremental remote replication.

(Source A --> DR1 B and Source A --> DR2 C) ==> (DR1 B --> DR2 C)

Use Case 4: Refresh or replace the source VNX using 1 to many replication.

Copy the data from VNX A (old) to VNX A (new) and synchronize VNX A (new) with VNX B. This allows site VNX A (old) to be synchronized to VNX A (new), remove VNX A (old), and increment establish between VNX A (new) and VNX B.

Requirement Attributes

- ◆ **Operational Interface** – CLI support only, no GUI support.
- ◆ **Interoperability** – Only V 6.0 configurations are supported.


Failover configuration rules

Note: Older versions of this document used the term *Utility Kit PowerPath* to refer to PowerPath SE.

PowerPath SE is for servers with *only one* HBA connected to the storage system. PowerPath SE supports only a *single path* to each SP, and *neither* supports load balancing. PowerPath versions 4.6 or later, which are MPIO-based, are for Windows Server 2003 servers. PowerPath version 4.5, which is not MPIO-based, is for Windows 2000 servers. PowerPath versions 4.5, 4.6, and higher support Fibre Channel and iSCSI configurations. Power versions 4.6 and higher replace PowerPath iSCSI, which is MPIO-based, for Windows Server 2003 servers in iSCSI configurations. PowerPath version 5.1.2 is the first PowerPath version to support Windows Server 2008 **or** Windows Server 2008 R2 servers. The term *PowerPath* refers to PowerPath with all its features, that is, support for servers with single or multiple HBAs with *single* or *multiple paths* to an SP and load balancing.

Supported configurations

Failover software is supported for any configuration with at least one path to each SP in a storage system. PowerPath or other multipath failover software is *required* for all multipath configurations; that is configurations with multiple paths from a server to an SP. [Figure 49 on page 141](#) through [Figure 52 on page 144](#) illustrate typical failover configurations:



Description: 1 HBA and 1 SP.

Supported: With or without PowerPath.

Failover ability: None, even with PowerPath installed


Description:

HBAs and 2 SPs.

Supported: With or without PowerPath, native failover, or DMP.

Failover ability: Requires PowerPath, PowerPath iSCSI, native failover, or DMP. PowerPath SE is *not* supported.

Figure 49 Direct connections between server and storage system



Description: 1 HBA and 2 SPs.

Supported: With or without PowerPath, native failover, server, or DMP.


Failover ability: Requires PowerPath SE, full PowerPath, native failover, or DMP.

Description: 2 HBAs and 2 SPs. Also applicable to similar configurations with more than 2 HBAs in the

Supported: With PowerPath, native failover, or DMP. For information on the type of failover supported (multipath or alternate path), see *Path rules for servers with failover software* (page 19).

Failover ability: Requires full PowerPath, native failover, or DMP.


Figure 50 SAN configurations (FC switch connections between server and storage system)



Description: 2 CNA and 2 SPs.

Supported: With PowerPath or native failover.

Failover ability: Requires full PowerPath, native failover, or DMP.



Description: 2 CNAs and 2 SPs. Also applicable to similar configurations with more than 2 CNAs in the server.

Supported: With PowerPath, native failover, or DMP. For information on the type of failover supported (multipath or alternate path), see *Path rules for servers with failover software* (page 19).


Failover ability Requires full PowerPath, native failover, or DMP.

Figure 51 FCoE switch connections between server and storage system

Description: 1 NIC or HBA and 2 SPs.

Supported: With or without PowerPath or native failover.

Failover ability Requires PowerPath SE, PowerPath or native failover.



Description: 2 NICs or HBAs and 2 SPs. Also applicable to similar configurations with more than 2 HBAs in the server.

Supported With PowerPath, PowerPath iSCSI, or native failover.

Failover ability: Requires PowerPath, PowerPath iSCSI or native failover.




Figure 52 Network configurations (iSCSI network connections between server and storage system)

Supported failover environments

IMPORTANT

If the server is part of a split-bus configuration, *none* of the following are supported:

- PowerPath
- PowerPath SE
- DMP with VNX for Block driver (VNX for Block driver is not applicable to DMP for AIX)

PowerPath

AIX, Linux, HP-UX, Net Ware, Solaris, VMware ESX, VMware ESXi, VMware ESX Server, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000

DMP

AIX, HP-UX, Linux, Solaris, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, or Windows 2000. For the supported revisions of VERITAS Volume Manager and any restrictions, see the *E-Lab Interoperability Navigator*. For DMP 4.x or earlier, a VNX for Block Array Specific Library (ASL) must be installed; ASLs are available from Symantec. DMP supplied with Veritas Volume Manager 5 does not required an ASL. If PowerPath is present and the array failover mode is either 1 or 4, PowerPath suppresses DMP control over LUNs. Power Path and DMP are not intended to be interoperable, in any other way.

Native failover

AIX - Multipath failover (MPIO)
 HP-UX - Multipath failover (MPIO), PVLinks
 Linux - Multipath failover (MPIO)
 Open VMS
 Solaris - Sun StorEdge Traffic Manager (SSTM)
 VMware ESX 4
 VMware ESXi 4
 VMware ESX 3 (ESX Server 3)
 VMware ESXi 3 (ESX Server 3i)
 VMware ESX Server 2.5
 Windows Server 2008 - Multipath failover (MPIO)
 Windows Server 2008 R2- Multipath failover (MPIO)

Asymmetric active active (ALUA) failover support

For information on which failover software for which operating systems support ALUA, see Knowledgebase emc99467.

Supported storage systems

For supported storage systems, refer to the *E-Lab Interoperability Navigator*.

HBA, CNA, or NIC types

Fibre Channel HBAs from the same HBA vendor, server bus type, and Fibre Channel link speed are required for each path to a storage-system Fibre Channel data port. For example, on a Solaris server you can connect either Emulex PCI or SBus HBAs, but *not* both types of HBAs, to the same storage system.

IMPORTANT

CNAs from the same CNA vendor and Ethernet link speed are required for each path to a storage-system FCoE data port.

Either NICs with the same Ethernet link speed or iSCSI HBAs from the same HBA vendor with the same Ethernet link speed are required for each path to a storage-system iSCSI data port. The same server *cannot* use both NICs and HBAs for paths to the same storage system.

Non-disruptive software installation (NDU) configuration rules

Supported storage systems

NDU is supported for all VNX storage systems.

Supported configurations

NDU supports any configurations with the failover software listed under *“Supported configurations” on page 140*.

When this document was published, there were certain restrictions regarding NDU operations for VNX storage systems connected to AIX server configurations.

For details and the latest status, review Knowledgebase emc67186 (AIX).

Reserved LUN pool

The reserved LUN pool works with replication software, such as SnapView, incremental SAN Copy, and MirrorView/A to store data or information required to complete a replication task. The reserved LUN pool consists of one or more private LUNs. A LUN becomes private when you add it to the reserved LUN pool. Since the LUNs in the reserved LUN pool are private LUNs, they cannot belong to storage groups and a server cannot perform I/O to them. [Table 46 on page 147](#) lists the maximum number of reserved LUNs for each storage system.

Before starting a replication task, the reserved LUN pool must contain at least one LUN for each source LUN that will participate in the task. You can add any available LUNs to the reserved LUN pool. Each storage system manages its own LUN pool space and assigns a separate reserved LUN (or multiple LUNs) to each source LUN.

All replication software that uses the reserved LUN pool shares the resources of the reserved LUN pool. For example, if you are running an incremental SAN Copy session on a LUN and a snapshot session on another LUN, the reserved LUN pool must contain at least two LUNs - one for each source LUN. If both sessions are running on the same source LUN, the sessions will share a reserved LUN.

Table 46 Number of reserved LUNs

Storage system	Maximum number of reserved LUNs
	VNX OE 05.31 or later
VNX5100	128
VNX5300	256
VNX5500	256
VNX5700	512
VNX7500	512

MirrorView/A and MirrorView/S configuration rules

A primary image can be a LUN or metaLUN (not a metaLUN component). A remote mirror connection between two storage systems can be used for MirrorView/A and/or MirrorView/S mirrors. However, a specific LUN can be involved in only one type of

mirror (asynchronous or synchronous, not both). [Table 58 on page 184](#) lists the MirrorView/A parameter limits for VNX series storage systems and [Table 59 on page 185](#) lists the MirrorView/S parameter limits for VNX series storage systems.

Note: MirrorView/A and MirrorView/S connections are not supported for storage-system FCoE data ports.

Supported storage systems for Fibre channel mirroring configurations

The following storage systems are supported for Fibre Channel mirroring configurations: VNX5100, VNX5200, VNX5300, VNX5400, VNX5500, VNX5600, VNX5700, VNX5800, VNX7500, VNX7600, VNX8000; CX4-120, CX4-240, CX4-480, CX4-960, CX3-10c, CX3-20, CX3-20c, CX3-20f, CX3-40, CX3-40c, CX3-40f, CX3-80, AX4-5, and AX4-5F8.

Supported storage systems for iSCSI mirroring configurations

The following storage systems are supported for iSCSI mirroring configurations: VNX5200, VNX5300, VNX5400, VNX5500, VNX5600, VNX5700, VNX5800, VNX7500, VNX7600, VNX8000; CX4-120, CX4-240, CX4-480, CX4-960, CX3-10c, CX3-20c, CX3-40c running FLARE 03.26.xxx.5.yyy or later.

Supported operating systems

All supported operating systems work with MirrorView/A and/or MirrorView/S.

SP port connections

- ◆ MirrorView/A and MirrorView/S SP ports can be connected via SP Fibre Channel
- ◆ MirrorView SP port connections for VNX storage systems are described in [“VNX series MirrorView ports” on page 94](#).

Fibre Channel MirrorView connections

SP A and SP B FC MirrorView ports (as described in [“VNX series MirrorView ports” on page 94](#)) on both the primary and secondary storage systems must run at the same Fibre Channel link speed.

You may use the same SP port for server data and MirrorView/A and/or MirrorView/S. Users should be cautious about sharing ports between MirrorView and server traffic when an IP distance connection is used. Sharing a MirrorView port with server traffic may cause a degradation in both replication and server application performance.

SP port Fibre Channel direct connections

- ◆ Direct Fibre Channel connections between the SP A MirrorView ports on the primary and secondary storage systems
- ◆ Direct Fibre Channel connections between the SP B MirrorView ports on the primary and secondary storage systems

SP port Fibre Channel switch connections

- ◆ Fibre Channel switch zone with SP A MirrorView ports on both the primary and secondary storage systems
- ◆ Fibre Channel switch zone with SP B MirrorView ports on both the primary and secondary storage systems

Note: If a server HBA port uses an SP MirrorView port for server I/O, it should be included in a separate zone.

SP port iSCSI connections

- ◆ iSCSI connections between the SP A MirrorView ports on the primary and secondary storage systems
- ◆ iSCSI connections between the SP B MirrorView ports on the primary and secondary storage systems

In an iSCSI mirroring configuration, you *cannot* create more than one path from a physical iSCSI port to another physical iSCSI port over different virtual ports.

Note: A storage system can have mirroring connections to a maximum of four other storage systems concurrently.

Additional considerations

If a primary image has multiple secondary images, the storage systems with the secondary images must have MirrorView connections. MirrorView connections are needed to promote one of the secondary images and still have all the secondary images as part of the mirror.

You must manage both the primary and secondary storage systems from a single Unisphere session running from either:

- ◆ A management station or storage system in the same domain as the primary and secondary storage systems.
- ◆ The local domain in a multi-domain environment that contains the domains with the primary and secondary storage systems

If any SnapView snapshot sessions and/or SAN Copy sessions are running on the same source LUN as a MirrorView/A mirror, these sessions share the same reserved LUN(s).

SAN Copy configuration rules

IMPORTANT

A storage system with SAN Copy software installed is called a *SAN Copy storage system*. A storage system with or without SAN Copy software that participates in SAN Copy sessions is called a *remote storage system*. An SP port on a SAN Copy storage system that participates in the SAN Copy session is called a *SAN Copy port*. Since a SAN Copy port functions like a server initiator, it can connect to only one storage group in a remote storage system. If MirrorView/A or MirrorView/S is installed, the MirrorView ports *cannot* be SAN Copy ports.

SAN Copy connections are not supported for storage-system FCoE data ports.

Remote storage system connections

A SAN Copy port can connect to a remote storage system in any of the following ways:

- ◆ Fibre Channel directly (see the SAN Copy release note for supported configurations)
- ◆ Through a Fibre Channel switch

- ◆ Through an FC-to-IP device
- ◆ Through a Fibre Channel switch to an FC-to-IP device
- ◆ iSCSI directly
- ◆ Through an iSCSI network

Supported operating systems for servers/storage systems

Table 47 lists the operating systems supported for servers connected to storage systems in SAN Copy session.

Table 47 Server operating systems for storage systems in SAN Copy

Connection between server and storage system	Operating system supported for server
Fibre Channel direct	HP-UX, Linux, Net Ware, Solaris, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000
Fibre Channel SAN	All operating systems - AIX, HP-UX, Linux, Net Ware, Solaris, VMware ESX 4, VMware ESXi 4, VMware ESX 3, VMware ESXi 3, VMware ESX Server 2.5, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, Windows 2000
iSCSI direct	AIX, HP-UX, Linux, Solaris, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, VMware ESX 4, VMware ESXi 4, VMware ESX 3, VMware ESXi 3, VMware ESXi, Solaris, Windows 2000
iSCSI network	AIX, HP-UX, Linux, Solaris, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003, VMware ESX 4, VMware ESXi 4, VMware ESX 3, VMware ESXi 3, Windows 2000

Support for storage systems in SAN Copy Fibre Channel configurations

Storage systems in a SAN Copy session

- ◆ All AX4-5 series, CX4 series, and CX3 series
- ◆ Supported non-VNX for Block storage systems in SAN Copy Fibre Channel configurations

See the *E-Lab Interoperability Navigator* for the current list of supported non-VNX storage systems.

Support for storage systems in SAN Copy iSCSI configurations

Storage systems in a SAN Copy session

CX4 series, CX3-10c, CX3-20c, CX3-40c, AX45SCi, AX4-5i

Operational restrictions

Operational restrictions apply when using SAN Copy with the Microsoft Cluster Server. Refer to the SAN Copy release notes on the EMC Online Support website for more information. For clusters other than the Microsoft Cluster Server, the following restrictions apply:

- ◆ If a SAN Copy source LUN is on a remote storage system, then either the LUN must be put in a state where it does not change during the entire copy operations (such as unmounting the LUN) or the copy should be done using a replica (snapshot or clone) as the source LUN.
- ◆ Various cluster software implements Persistent Reservations on the LUNs, which could prevent SAN Copy from being able to access one or more LUNs involved in a session. (See Knowledgebase emc87610, which, although it is written for Symmetrix systems, notes which clustering software uses Persistent Reservations). For this reason, EMC recommends that you:
 - Do not add the LUNs on the remote storage system participating in a SAN Copy session to Storage for the cluster.
 - Before creating or starting a SAN Copy session whose source or destination LUNs are on a remote storage system, take those LUNs offline. If you cannot take the LUNs offline, then create a snapshot of the LUNs on the remote storage system participating in a SAN Copy session. Add these snapshots to the storage group for the cluster.

Logical unit ownership

A logical unit for a VNX system is either a LUN or a metaLUN (not a component of a metaLUN). For a Symmetrix system, a logical unit is a volume. Either the source logical unit or all the destination logical units in a SAN Copy session must be in the SAN Copy storage system that owns the session; that is, on the storage system from which you started the session. If the source logical unit resides in the SAN Copy storage system, you can copy data to one or more destination logical units in one or more remote storage systems. [Table 60 on page 185](#) lists the SAN Copy parameters and storage system limits.

For incremental SAN Copy sessions, the source logical unit must reside in the SAN Copy storage system. This means, for example, that you *cannot* use a Symmetrix system Volume as a source logical unit for an incremental SAN Copy session.

SP port usage

In a VNX SAN Copy storage system or remote storage system, any SP port can participate in a SAN Copy session, except for the ports used for mirroring if MirrorView/A or MirrorView/S is installed. Since a SAN Copy port acts like an initiator, it can connect to only one storage group on a remote storage system. As a result, EMC recommends that all the LUNs participating in SAN Copy sessions in a remote storage system be in the same storage group. When a server initiator (HBA) is using a SAN Copy port to transfer I/O, that port acts as a target to the server initiator.

In an iSCSI SAN Copy configuration, you *cannot* create more than one path from a physical iSCSI port to another physical iSCSI port over different virtual ports.

Zoning requirements

Because the SAN Copy port acts like an initiator, SAN Copy is an exception to EMC's single-initiator, multi-target zoning recommendation. A zone containing an HBA port can contain *only one* SAN Copy SP port.

IMPORTANT

Since a SAN Copy port acts like an initiator, it logs in as an initiator on any other SP, FA, or FC ports to which it is zoned. Each time a server initiator registers with the SAN Copy port or the SAN Copy port registers with another SP port, the SAN copy port consumes a login resource. If you consume all login resources for the SAN Copy port, the SAN Copy session may not successfully complete to all destination logical units.

For high availability when you use SAN Copy, you should configure a minimum of two zones. Each zone should include one port from a different SP in the SAN Copy storage system and one port from each SP, channel adapter, or controller in the remote storage system. For more information on SAN Copy zoning, refer to the *EMC SAN Copy for Navisphere Administrator's Guide*.

Additional considerations

If any MirrorView/A and/or SnapView snapshot sessions are running on the same source LUN as a SAN Copy session, these sessions share the same reserved LUNs.

For information on using SAN Copy with Microsoft Cluster Server (MSCS), refer to the SAN Copy release notes on the EMC Online Support website.

LUN usage with replication applications

Replication applications used with LUNs are SnapView, SAN Copy, MirrorView/A, and MirrorView/S. [Table 48](#) lists the LUN types with replication features and their potential use.

Table 48 LUN usage with replication applications (page 1 of 2)

LUN type	Potential usages				
	Make a snapshot of it?	Make a clone (BCV) of it?	Use as source for MirrorView?	Use as source for <i>full</i> SAN Copy session?	Use as source for <i>incremental</i> SAN Copy session?
LUN or metaLUN not yet in use with any replication application ^a	Yes	Yes	Yes Can use with MirrorView/A or MirrorView/S but not both at once	Yes CAUTION – Source LUN should not change during the copy	Yes
Thin LUN	Yes	Yes	Yes	Yes	Yes
Snapshot	Yes Up to 8 snapshots of a snapshot	No	No Cannot mirror a snapshot	Yes	No Incremental sessions use snapshots internally

Table 48 LUN usage with replication applications (page 2 of 2)

LUN type	Potential usages				
	Make a snapshot of it?	Make a clone (BCV) of it?	Use as source for MirrorView?	Use as source for <i>full</i> SAN Copy session?	Use as source for <i>incremental</i> SAN Copy session?
Clone (BCV)	Yes Up to 256 VNX snapshots per clone	No Cannot clone a clone	No Cannot mirror a clone	Yes CAUTION – Source LUN should not change during the copy	Yes
MirrorView source LUN	Yes	Yes ^b	No Cannot mirror a mirror	Yes CAUTION – Source LUN should not change during the copy	Yes
MirrorView destination LUN	Yes	Yes ^b	No Cannot mirror a mirror	No Create snapshot and use as the source for SAN Copy session	Yes

a. Includes any LUN used as a target for either a *full* or *incremental* SAN Copy session. A LUN in a RAID group with power savings enabled cannot be used with replication applications.

b. Only if the source LUN is either in a CX3 series storage system running FLARE 03.24.xxx.5.yyy or later.

VNX Snapshots

VNX Snapshots provide a method of creating different types of copies of source LUNs on a VNX array. The method is commonly referred to as snapshots. Unlike a clone, which is an actual copy of a LUN, a snapshot is a virtual point-in-time copy of a LUN. A snapshot session can be started at any time. The snapshot session keeps track of how the source LUN looks at a particular point in time. VNX Snapshot includes the ability to roll back or restore a snapshot session.

VNX Snapshots (or write-in-place pointer-based snapshots) support Block LUNs only and require pool-based LUNs (thick LUNs and thin LUNs). Up to 256 writeable snapshots per LUN are supported.

Snapshots allocate their space from the storage pool, not from thin or thick source LUNs. Snapshots can optionally be migrated to pool LUNs.

Thin or Thick LUNs

VNX Snapshots can be applied to either thin or thick LUNs. A thick LUN contains 8KB of granular mapping overhead (a non-sequential index map) when the first snapshot is created. This overhead can reduce system performance since:

- ◆ Subsequent writes have 8KB of granular mapping overhead, so performance will be comparable to a thin LUN.
- ◆ Reads to changed original data have their performance comparable to a thin LUN.
- ◆ Reads to unchanged original data have the original thick LUN performance.

After its last snapshot is removed, the thick LUN's performance advantage is restored. The LUN index is converted back to direct mapping (sequentially ordered data within each 1 GB slice).

Snapshots of snapshots

Snapshots of snapshots (or hierarchical snapshots) are limited by the number of snapshots per LUN limit. Up to 8 levels of hierarchical snapshots are supported.

Auto-delete

VNX Snapshots support configurable snapshot auto-delete (or harvesting). By default, auto-delete removes the oldest snapshots when pool is nearly full. When the pool is full, auto-delete removes all the snapshots. The user can permanently protect selected snapshots from auto-deletion.

SnapView snapshots and VNX snapshots

SnapView snapshots and VNX snapshots are allowed on the same LUN. This is useful for pre-existing SnapView snapshots on pool LUNs, and/or existing scripts using pre-existing SnapView snapshots, which must not be broken.

Clone support

VNX Snapshots support clones -- a separate local protection technology. Clones are limited to 8 per source LUN and up to 256 VNX snapshots per clone are supported. Each clone can either be a thick or thin LUN. A clone source can also be thick or thin.

VNX Snapshot management

VNX Snapshots are managed in Unisphere GUI, from the Protection section (where SnapView, SnapSure, and Clone management options exist). VNX Snapshots can also be managed directly via Unisphere CLI.

VAAI support for NFS

The vStorage Application Program Interface (API) for Array Integration (VAAI) for NFS allows a NAS vendor to develop a VMWare ESX server vStorage plug in that creates and manages virtual disks as a set of files on a NAS device. In a vStorage environment, the system administrator creates one golden copy of the virtual disk image, stores it as a file on a NAS server, and then creates thousands of writeable-snaps for virtual desktop deployment.

The vStorage API specifies various functionality that a NAS device must support: fast clone, full clone, space reservation, extended file stat, and snap of snap. The vStorage API also defines the primitives between VMWare virtual disk library and the ESX vStorage plug in. The primitives include Init_Session, Execute_primitive, Query and End_session. The protocol between the ESX plug in and the NAS device is to be determined by the NAS vendor.

The overall architecture consists of multiple parts:

- ◆ ESX vStorage plug in
- ◆ Data Mover vStorage enabler
- ◆ Snap of snap of VMDK files

The ESX vStorage plug in works as a dynamic linked library inside the VMWare ESX host. It interacts with the ESX virtual disk library via the defined primitives. The plug in translates the primitive into an RPC, communicates with the Data Mover vStorage enabler, obtains the result, and replies to the ESX virtual disk library.

The Data Mover vStorage enabler is a protocol interpreter as well as the entry point to invoke file services in order to service the requests. The Data Mover vStorage enabler also keeps the necessary state such as sessions and tasks that are relevant to the requests from the ESX vStorage plug in.

For fast clone, the design and implementation is based on file versioning technology. A new file can be created as a version file and an existing file can be converted to a version file. A fast clone is essentially a writeable snap of the working file. Snap will be created instantaneously.

For the creation of a full clone, buffered I/O interfaces are used to read data from the source file and write data to the destination file. For resource consumption, throttling, and performance optimization, the copy process can be multi-threaded with memory consumption capped at a pre-determined upper limit (currently proposed as 64MB).

For space reservation, leveraging the persistent reservation (PBR) functionality would fulfill the requirement. By default, a file created using NFS protocol is a sparse file. A sparse file can be converted to a dense file via the NFS API extension. The latter would satisfy the requirement of a reserve space primitive on a given file.

Snap of snap of VMDK files on NFS file systems is supported at a minimum of one level of depth (source to snap to snap). With this functionality, VMware View can leverage VAAI for NFS on a VNX system rather than taking the snaps at host/ESX level. This allows VAAI for NFS to take snaps of snaps of desktop images.

Figure 53 depicts the high-level architecture.




Figure 53 VAAI High Level Architecture Diagram

The ESX vStorage plug in is a user space program inside the VMWare ESX host, interacting with the virtual disk library via VAAI primitives. Upon receiving the request from the virtual disk library, the ESX vStorage plug in locates the proper Data Mover and file handle, and then issues a NFS v3 call to perform the desired operation.

The VAAI Enabler interprets the command and executes the operation. In addition, the VAAI Enabler also monitors the progress of the long running command and informs the ESX vStorage plug in of the operation status via an NFS v3 extension.

VASA support for VMware

The vStorage API for Storage Awareness (VASA) is a proprietary standard operating procedure-based Web interface used by VMware clients (typically VMware Server/Infrastructure Administrators). It is aimed at the VMware Administrator using the VMware vCenter (VMware component that acts as the client).

VMware uses the VASA Provider Interface to request basic information about the array and the storage devices it exposes to the virtual environment. This is used to facilitate day-to-day monitoring and troubleshooting through VASA.

Note: VASA is a reporting interface, there are no VASA APIs for active management.

The Solutions Enabler VASA Provider (SE VP) supports the following arrays:

- ◆ VNX arrays
- ◆ Installed base support for legacy arrays (CX4 and NS block)

VAAI support for block

The vStorage API for Array Integration (VAAI) for block allows the integration of block VAAI support into the core storage stack and support for thin provisioning.

VAAI is composed of several primitives including: Block Zeroing, Cloning File Blocks, and Hardware Assisted Locking. VMWare supports these primitives by T10 standard commands integrated into the core storage stack. This allows easier standard based integration with vendors and better performance by off loading to the hardware primitives.

The T10 standard commands integrated into the core storage stack include:

- ◆ Extended Copy Results - copy data from a set of logical units to the same or another set of logical units.
- ◆ Receive Copy Results - displays the results of the Extended Copy command.
- ◆ Write Same (16) - a 16-byte command that writes a block of data multiple times to consecutive blocks starting at the logical block address (LBA).
- ◆ Vital Product Data page 0x83 (VPD 83) - the device identification page.
- ◆ Compare and Write (not fully integrated)
- ◆ UNMAP - for thinly provisioned LUNs, used to free physical space on a LUN when a VMFS file is deleted.
- ◆ Compare and Swap - instruct the target to accept data to compare and conditionally store the data to the given LUN. This command operates as a WRITE (16) command except with respect to preventing any other CAS updates to the target blocks on the LUN while this command is underway. The write is performed only if the specified compare with existing data on the LUN is successful.

- ◆ Vital Product Data page B0 (block limits page) - used to determine if a LUN supports the UNMAP command and returns data about the arrays block support.

CHAPTER 3

VNX R31-R32 SW Configuration Rules

This chapter covers configuration rules for storage arrays running VNX R31 and R32. This VNX software runs only on VNX5100, VNX5300, VNX5500, VNX5700, VNX7500, VNX VG2, and VNX VG8. Topics include:

◆ Fibre Channel storage area network (SAN) configuration rules	164
◆ iSCSI network configuration rules	164
◆ Read/write cache configurations	165
◆ FAST Cache for VNX storage systems	168
◆ Compression for VNX series storage systems	169
◆ Storage-system parameters for VNX series storage systems	170
◆ Server software rules	175
◆ SnapView configuration rules	178
◆ MirrorView/A and MirrorView/S limits	184
◆ SAN Copy configuration limits	185
◆ VNX RecoverPoint write Splitter usage	186

Fibre Channel storage area network (SAN) configuration rules

The number of initiator records for the VNX series with VNX Operating Environment (OE) 5.31 and 5.32 for FC are shown in [Table 49](#).

Table 49 Initiator records for Fibre Channel

Storage system	Initiator records							
	Per FC port		Per FCoE port		Per SP		Per storage system	
	R31	R32	R31	R32	R31	R32	R31	R32
VNX5100 ^a	255	256	0	0	256	256	512	0
VNX5300	255	256	256	256	1024	1024	2048	2048
VNX5500	255	256	256	256	2048	2048	4096	4096
VNX5700	256	256	256	256	2048	2048	4096	4096
VNX7500 (24G)	512	512	256	512	4096	4096	8192	8192
VNX7500 (48G)	N/A	512	N/A	512	N/A	4096	N/A	8192

a. The VNX5100 ports available are all on the mezzanine card. This model does not support additional ports on I/O modules. See [“VNX5100” on page 31](#).

iSCSI network configuration rules

The number of initiator records for the VNX series with VNX OE for Block 05.32 for iSCSI are shown in [Table 50](#). The VNX5100 does not support iSCSI connections.

Table 50 Initiator records for iSCSI

Storage system	Initiator records	
	Per 1 Gb/s Ethernet iSCSI port	Per 10 Gb/s Ethernet iSCSI port
VNX5300	256	256
VNX5500	512	512
VNX5700	1024	1024
VNX7500 (24G)	1024	1024
VNX7500 (48G)	1024	1024

Each path used in a MirrorView or SAN Copy relationship between two storage system counts as an initiator for *both* storage systems.

Read/write cache configurations

The VNX-series has larger caches than previous EMC mid-range storage systems. These caches are highly configurable. The allocation of read and write cache can be tuned to achieve optimal performance for individual workloads.

The amount of available SP memory depends on the VNX model. This memory is shared by the storage systems operating environment, enabled VNX features, installed applications, and the read/write cache. An example of an enabled feature includes FAST Cache. An example of an installed application is RecoverPoint™. The operating environment, features, and installed applications use SP memory capacity called *system memory*. A partial list of the features that reduce system memory are:

- ◆ Installed Applications: Onboard, Management, RecoverPoint
- ◆ Layered Applications: SnapView, MirrorView
- ◆ Features: Virtual Provisioning, FAST Cache

The storage systems memory is a resource that must be managed. The capacity of the read/write cache determines a large part of the storage systems performance. The number of enabled features and installed applications reduces the amount of memory available for the storage systems read/write cache. Be aware of the effect a

reduced capacity read/write cache may have on overall performance. [Table 51](#) shows the maximum cache allocations for the VNX Operating Environment in block and file configurations.

Table 51 Maximum cache allocations

	VNX5100		VNX5300		VNX5500		VNX5700		VNX7500 (24)		VNX7500 (48)
	R31	R32	R31	R32	R31	R32	R31	R32	R31	R32	R32
Maximum Storage Processor Memory (GB)	4	4	8	8	12	12	18	18	24	24	48 (R32)
Maximum SP Cache (GB)	1.442	.950	3.385	3.997	6.294	6.488	11.3	9.706	12.8	13.45	37.121
Maximum Storage Processor Write Cache with Data Services Installed (MB)	500	650	2162	2782	3396	4738	7782	6581	6451	10600	16600
CBFS Cache [Data Services] (MB)	120	123	240	293	360	416	540	651	720	1106	10322
Cache Size for Customer Data (with Data Services) (MB)	620	773	2402	3075	3756	5154	8322	7232	7171	11706	26992

To use read or write caching, you must set up read or write caching for the storage system and for the LUNs that will use caching. Setting up read or write caching for a storage system requires that you assign SP memory to the cache and set the storage-system cache properties. Set up read or write caching for a LUN when you create it or after it is bound by setting the LUN's cache properties. The default settings for the LUNs cache created by Flash drives is selecting both read and write cache ON.

Use cases for using the default settings are:

- ◆ Read and write cache is extended to LUNs created on Flash drives without any adverse effect on any other LUNs
- ◆ Write IOPS are significantly improved in newer revisions of VNX OE for Block
- ◆ Most arrays are deployed with adequate write cache flushing overhead

Allocating read and write cache recommendations

Allocate system memory first, then allocate *the rest of* the memory to the read/write cache.

For the majority of workloads, the default ratio allocation of read/write cache should be left unchanged. Note that enabling features will maintain this ratio, although the overall capacity of the read/write cache may be reduced. The capacity of the storage systems read/write cache is the remainder of the SP memory *after* system memory has been allocated. In addition, note that enabling some features causes the read/write cache to be temporarily disabled as the SP memory is re-allocated.

It is advisable to have a minimum of 200 MB, with a maximum of 20-percent of the available memory allocated to read cache. However, it is always important to maximize the storage systems write cache size. If performance is suffering from a small write cache, lower the write cache watermarks to free-up cache pages more quickly to make-up for a small write cache. Allocate all of the available memory to be used as write cache, which means operate without read cache.

Not having a read cache allocation will adversely affect the storage systems sequential read performance, but having a larger write cache will have a more positive effect on overall system performance.

Note: A write cache allocation applies to both SPs; it is mirrored. Read cache is not mirrored.

Because of write cache mirroring, each of the two SPs receives the same allocation taken from memory. For example, if you allocate 1 GB to write cache, 2 GB of the available system memory is consumed.

When you allocate memory for read cache, you are only allocating it on one SP. SPs always have the same amount of write cache, but they may have different amounts of read cache. To use the available SP memory most efficiently, allocate the same amount of read cache to both SPs.

FAST Cache for VNX storage systems

Storage systems with VNX OE for Block 05.32 or later support an optional FAST Cache consisting of a storage pool of Flash disks that are configured to function as a FAST Cache. This cache provides low latency and high I/O performance without requiring a large number of Flash disks. It is also expandable while I/O to and from the storage system is occurring.

The FAST Cache is based on the locality of reference of the data set requested. A data set with high locality of reference that is frequently accessed is a good candidate for promotion to the FAST Cache. By promoting the data set to the FAST Cache, the storage system services any subsequent requests for this data faster from the Flash disks that make up the FAST Cache; thus, reducing the load on the disks in the LUNs that contain the data (the underlying disks). Applications such as file and OLTP (online transaction processing) have data sets that can benefit from the FAST Cache.

Note: The use of FAST cache (actually the enabler being installed) reduces the system’s DRAM cache.

The FAST Cache consists of one or more pairs of mirrored disks (RAID 1 type) and provides both read and write caching. For a read caching, the FAST Cache copies data off the disks being accessed into the FAST Cache. For protected-write applications, the FAST Cache effectively buffers the data waiting to be written to disk. In both cases, the workload is off-loaded from slow rotating disks to the faster Flash (SSD) disks in the FAST Cache. The performance boost provided by the FAST Cache varies with the workload and the cache size. The FAST Cache can be enabled for specific LUNs in RAID groups or all the LUNs in one or more pools. The FAST Cache should be disabled for Write Intent Log (WIL) LUNs or clone private LUNs. Enabling the FAST Cache for these LUNs is a misallocation of the FAST Cache and may reduce the effectiveness of the FAST Cache for other LUNs. [Table 52](#) lists the storage system type, Flash disk capacity, number of Flash disks in FAST Cache, and the maximum FAST Cache capacity.

Table 52 FAST Cache and disk configurations

Storage system	Flash (SSD) disk capacity	Number of Flash (SSD) disks in FAST Cache	Max FAST Cache capacity
VNX5100	100 GB	2	100 GB
	200 GB	N/A	N/A

Table 52 FAST Cache and disk configurations (continued)

VNX5300	100 GB	10	500 GB
	200 GB	4	400 GB
VNX5500	100 GB	20	1000 GB
	200 GB	10	1000 GB
VNX5700	100 GB	30	1500 GB
	200 GB	14	1400 GB
VNX7500	100 GB	42	2100 GB
	200 GB	20	2000 GB

VNX5100 rules for enabling FAST Cache or Thin Provisioning

Safeguards are implemented as part of software installation process to allow either FAST Cache (to boost performance) or Thin Provisioning Enabler (to optimize capacity utilization) to be installed on VNX5100 systems.

This allows the user to have a choice between these two features without risking adverse impact to system performance. It also eliminates the possibility of a user from accidentally installing both the enablers on the same system.

Note: FAST Virtual Provisioning is not available for the VNX5100.

Compression for VNX series storage systems

VNX series storage systems with VNX OE for Block 05.32 and both the Compression and Thin Provisioning enablers provide the option of compressing the data on the LUN to free up storage space. Compression analyzes the data on a disk and applies algorithms that reduce the size of repetitive sequences of bits that are inherent in some types of files. Unlike de-duplication, compression does not depend on multiple copies of the same data. The storage system performs compression operations in the background while continuing to service the host I/O.

Compression of a RAID group or thick LUN converts the LUN to a thin LUN. The compressed LUN remains a thin LUN even after you decompress it, but the LUN's host-visible size is equal to the size that you specified when you created the LUN. You

cannot compress a newly created thick LUN without data on it. For these reasons, the recommended procedure for creating a compressed LUN is to create the LUN as a thin LUN, and once created, turn on compression for the LUN.

Storage-system parameters for VNX series storage systems

Table 53 list the storage-system parameters for VNX series storage systems with VNX OE for Block 05.31 - 32 which include:

- ◆ Basic system parameters
- ◆ RAID group and RAID group LUN (traditional) parameters
- ◆ Pool and pool LUN parameters
- ◆ Storage group parameters
- ◆ Domains
- ◆ Compression parameters
- ◆ Unisphere Quality of Service Manager parameters

Table 53 Storage system parameters (page 1 of 5)

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Basic system parameters					
Max FC ports per SP (FE ports) Mezzanine	4	4	4	0	0
Max FC ports per SP (FE ports) SLIC Ports	0	4	4	12	16
Max I/O modules per SP	0	2	2	5	5

Table 53 Storage system parameters (page 2 of 5)

Max 1Gbps iSCSI ports per SP (FE only)	0	4	8	12	12
Max 10Gbps iSCSI ports per SP (FE only)	0	4	4	6	8
Max FCoE ports per SP (FE only)	0	4	4	6	8
Initiator records per FC port (max)	255	255	255	256	512
Initiator records per 1 Gb/s Ethernet iSCSI port (max)	0	256	512	1024	1024
Initiator records per 10 Gb/s Ethernet iSCSI port (max)	0	256	512	1024	1024
Initiator records per 10GBase-T iSCSI/IP port (max)	0	256	512	1024	1024
Max initiators/SAS port	0	128	128	128	128
Initiator records per SP (max)	256	1024	2048	2048	4096
Initiator records per storage system (max)	512	2048	4096	4096	8192
Front-end data ports per SP (max)	4	12	12	16	16 ^a
Back-end data ports per SP	2	2	2	4	8
Disks (max)	75	125	250	500	1000
Disks (min)	4	4	4	4	4

Table 53 Storage system parameters (page 3 of 5)

Hosts per array	128	512	1024	1024	2048 (24) 4096 (48)
Virtual machines per array	128	640	1280	2560	5120
LDAP services	2	2	2	2	2
Virtual Centers	10	20	10	10	10
RAID group and RAID group LUN (traditional) parameters					
Disks per RAID group (max)	16	16	16	16	16
RAID groups (max)	75	125	250	500	1000
User-visible classic LUNs (max)	512	2048	4096	4096	8192
LUNs per RAID group (max)	256	256	256	256	256
metaLUNs (max)	256	512	512	1024	2048
LUN or metaLUN size (max) ^b	See note b .				
LUNs per metaLUN component (max)	32	32	32	32	32
MetaLUN volumes (max)	256	512	512	1024	2048
Components per metaLUN (max)	8	8	8	16	16
LUNs in a MetaLUN (max)	256	256	256	512	512
Max total consumables	512	2048	4096	4096	8192; 16384 for 7500-48 (R32)

Table 53 Storage system parameters (page 4 of 5)

RAID types ^c	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk
Pool and pool LUN parameters					
Disks in all pools (max)	71	121	246	496	996
Disks per pool (max)	71	121	246	496	996
Pools (max)	10	20	40	40	60
LUNs per pool (max)	512	512	1024	2048	2048
Pool LUNs per storage system (max)	512	512	1024	2048	2048
Pool LUN size (max)	16 TB	16 TB	16 TB	16 TB	16 TB
RAID types	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0
Storage group parameters					
Storage groups (max)	128	256	512	1024	2048
LUNs per storage group (max)	256	512	512 (R30); 1024 (R32)	1024 (R30); 2048 (R32)	1024 (R30); 4096 (R32)
Max User Visible LUNs	896 (31) 1408 (32)	2816 (31) 3328 (32)	4864 (31) 5888 (32)	5376 (31) 7424 (32)	10496 (31) 12544 (32)
Max User Visible FLARE LUNs	512	2048	4096	4096	8192
Max LUNs per RAID group	256	256	256	256	256
Max LUNs per storage group	256	512	512 (31) 1024	1024 (31) 2048 (32)	1024 (31) 4096 (32)

Table 53 Storage system parameters (page 5 of 5)

Max number of disks that can be added to a pool at one time	40	40	80	120	180
Domains					
Max systems per domain	100	100	100	100	100
Max systems per multi-domain	100	100	100	100	100
Max domains per multi-domain	10	10	10	10	10
Compression parameters					
Concurrent compression operations	N/A	5	5 (R31); 10 (R32)	8 (R31); 16 (R32)	10 (R31); 20 (R32)
Unisphere Quality of Service Manager parameters					
Policies per storage system (max)	10	10	10	10	10
I/O classes per storage system (max)	64	64	64	64	64
LUNs per I/O class (max)	128	128	128	128	128
I/O classes per cruise control policy (max)	2	2	2	2	2
I/O classes per limit or queue depth policy (max)	32	32	32	32	32
Scheduled events (max)	10	10	10	10	10

a. In an 8 bus VNX7500, the max number of FE ports would be 12.

- b. Max Pool LUN is 16 TB. Max RG LUN can be larger than 16 TB. Larger LUNs can be configured provided host operating system can consume those LUNs.
- c. For best practices, use RAID 6 for less efficient drives and RAID 5 or RAID 1/0 for more efficient drives.

MetaLUNs

Configuration rules for metaLUNs are:

- ◆ A RAID 6 metaLUN can include only RAID 6 LUNs. For other types of metaLUNs, all metaLUN components must be either redundant (RAID 1, RAID 1/0, RAID 3, or RAID 5) or non-redundant (RAID 0 or Disk).
- ◆ All components in a metaLUN stripe must have the same size and RAID type.
- ◆ All disks in a metaLUN component must have the same type of disks.

Server software rules

Operating system

All servers attached to the storage system must run one of the following operating systems listed in [Table 54](#)

Table 54 Server operating systems (page 1 of 2)

Operating system	Storage systems supported ^a
	VNX Series
AIX ^b	Yes
HP-UX ^b	Yes
Linux	Yes
Open VMS ^b	No
Solaris ^b	Yes
VMware ESX 4, ESXi 4	Yes
Windows Server 2008	Yes

Table 54 Server operating systems (page 2 of 2)

Operating system	Storage systems supported ^a
	VNX Series
Windows Server 2008 Hyper-V	Yes
Windows Server 2008 R2	Yes
Windows Server 2003	Yes

a. Only limited configurations are supported for some storage systems. For more information, see *E-Lab Interoperability Navigator*.

b. Not currently supported for servers connected to storage-system FCoE data ports.

Unisphere Host Agent and Unisphere Server Utility

EMC recommends that you install the Unisphere Server Utility (not available for Open VMS) on all servers. Depending on your application needs, you can also install the host agent on each server connected to a storage system to:

- ◆ Monitor storage-system events and notify personnel by e-mail, page, or modem when any designated event occurs.
- ◆ Retrieve the LUN world wide name (WWN) and capacity information from Symmetrix® storage systems.
- ◆ Register the server's NICs or iSCSI HBAs with the storage system.

Alternatively, you can use the Unisphere Server Utility to register the server's NICs or iSCSI HBAs with the storage system. You can also use the Navisphere CLI to register the server's HBAs or NICs with the storage system. The host agent also sends operating-system device (LUN) mapping information to the storage system, which the server utility and the CLI *cannot* do.

Clustered server configuration rules

Clustering software is required whenever multiple servers running the same operating system are connected to the same storage system (non-storage groups) or to the same storage group. The servers must be part of the same cluster.

Clusters with VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500 storage systems

A cluster with VNX5100, VNX5300, VNX5500, VNX5700, or VNX7500 storage systems must use either iSCSI or Fibre Channel front-end data ports and *not* both types of ports.

Number of nodes

For information on the supported number of nodes and the minimum Microsoft Initiator version, see the *E-Lab Interoperability Navigator*.

Supported operating systems

All operating systems are supported for the storage system and ESX Server virtual machines. Some operating system clusters are not supported with iSCSI attaches. Refer to the *E-Lab Interoperability Navigator* for more information on supported operating systems.

The number of HBA, CNA, or NIC ports depends on the operating system and the failover software on the server. For information on the failover software supported for a particular operating system, see the [“Failover configuration rules” on page 140](#).

For PowerPath - two or more per server, except for Linux, Solaris, Net Ware, Windows Server 2008, Windows Server 2008 R2, Windows Server 2003 and Windows 2000, which can have one or more per server.

HBA, CNAs, or NIC types

Fibre Channel HBAs from the same HBA vendor, server bus type, and Fibre Channel link speed are required for each path to a storage-system Fibre Channel data port. CNAs from the same CNA vendor and Ethernet link speed are required for each path to a storage-system FCoE data port. Either NICs with the same Ethernet link speed or iSCSI HBAs from the same HBA vendor with the same Ethernet link speed are required for each path to a storage-system iSCSI data port. The same server *cannot* use both NICs and HBAs for paths to iSCSI storage systems. The server cannot connect to both storage-system Fibre Channel or FCoE data ports and storage-system iSCSI data ports, except as described in [“Fibre Channel, FCoE, and iSCSI storage-system connections for a single server” on page 105](#).

Paths to SP

See [“Read/write cache configurations” on page 165](#).

Software requirements

Cluster Software

See the *E-Lab Interoperability Navigator* for the cluster software supported for each operating system. Note that an MSCS cluster running Windows PowerPath 4.3.1 or later can connect to VNX, VNX for Block, and Symmetrix storage systems.

Failover Software

PowerPath where available. PVLinks for HP-UX. Multipath failover (MPIO) for ESX Server, ESX, ESXi, HP-UX, Linux, Solaris, AIX, Windows Server 2008 R2, and Windows Server 2008.

Notes

For information on cluster support over long distances, refer to the *E-Lab Interoperability Navigator*.

SnapView configuration rules

Note: *LUN* refers to *LUN* or *metaLUN* (not a component of a metaLUN).

IMPORTANT

Do not allocate SnapView reserved LUNs or clones from the same RAID Group as the Source LUNs that will be replicated. Doing so impacts performance severely.

Supported storage systems

The following storage systems are supported for SnapView: VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500.

Table 55 SnapView Limits

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Max Snapshots/Source LUN	8	8	8	8	8
Max Sessions/Source LUN	8	8	8	8	8
Max Snapshot Source LUNs	128 (31) 256 (32)	256	256	512	512

Table 55 SnapView Limits

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Max Snapshot LUNs (per array)	256 (31) 512 (32)	512	512	1024	2048
Max SnapShot Sessions (per array)	256	512	512	1024	2048
Max Snapshot Source LUNs in a Consistent Start	32	32	32	64	64

Source LUNs for clones

Any source LUN or thin LUN can be cloned, *except for*:

- ◆ Hot spare LUNs
- ◆ Clone LUNs (LUNs participating in any clone group as either a source LUN or a clone LUN)
- ◆ Snapshot LUNs
- ◆ Private LUNs (LUNs reserved as clone private LUNs in a reserved LUN pool or in a write intent log)
- ◆ A LUN in the process of conversion

[Table 56 on page 179](#) lists the SnapView clone parameters and their limits for VNX storage systems.

Table 56 SnapView supported limits for VNX series storage systems

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Clones (clone images)^a					
Clones per storage system	128 (R31); 256 (R32)	256	512 (R31); 1024 (R32)	1024 (R31); 2,048 (R32)	2048
Clone source LUNs	64 (31) 128	128	256 (31) 512 (32)	512 (31) 1024 (32)	1024
Clones per consistent fracture	32	32	32	64	64

Table 56 SnapView supported limits for VNX series storage systems (continued)

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Clone groups					
Clones in a Clone group	8	8	8	8	8
Clone private LUNs^{b, c}					
Per storage system (required)	2	2	2	2	2
Clone source LUNs					
Per storage system	64	128	256	512	1024

- a. The source LUN and MirrorView/S primary and secondary images no longer count towards the clone image limit.
- b. A thin LUN cannot be a clone private LUN.
- c. The minimum size for a clone private LUN is 1 GB.

Table 57 lists the SnapView snapshot parameters and their limits for VNX storage systems.

Table 57 Maximum SnapView snapshot limits for VNX series storage systems

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Snapshots^a					
Per storage system	256	512	512	1024	2048
Per source LUN	8	8	8	8	8
SnapView sessions^a					
Per source LUN	8	8	8	8	8
Per storage system	256	512	512	1024	2048
Source LUNs in consistent session	32	32	32	64	64
Reserved LUNs^b					
Per storage system	128	256	256	512	512
Source LUNs					
Per storage system	128 (R30); 256 (R32)	256	256	512	512

a. The limits for snapshots and sessions include SnapView snapshots or SnapView sessions, as well as reserved snapshot or reserved sessions used in other applications, such as and MirrorView/A and SAN Copy (incremental sessions).

b. A thin LUN cannot be a reserved LUN.

Note: You can create a snapshot of a clone, but you cannot create a clone of a snapshot. You can create a snapshot of a mirrored LUN, but you cannot create a clone of a mirrored LUN.

Supported operating systems

All supported operating systems work with SnapView. The server that uses the source LUN and the server that uses the replica (snapshot or clone) must run the same operating system.

Number of servers

Two or more, where one or more servers have access to the source LUN and one or more servers have access to the replica (snapshot or clone). EMC supports placing a clone or snapshot in the same storage group as its source LUN *only for a server running*

- ◆ Windows Server 2008, Windows Server 2008 R2, Windows Server 2008 Hyper-V, Windows Server 2003, or Windows 2000.
- ◆ VMware ESX 4, VMware ESXi 4, VMware ESX 3, VMware ESXi 3, VMware ESX Server 2.5 with the clone or snapshot is presented to a different virtual machine than its source LUN.
- ◆ Other operating systems, if you use RM/Local, RM/SE, or PowerPath Volume Manager to put the clone or snapshot in the storage group.

These operating systems and software products provide “same server access” to the clone, snapshot, and the source LUN. For information on using these operating systems or software products, refer to the documentation for the operating system or product.

Clustered servers

You can configure a clustered server to access a source LUN, but not to access both the source LUN and its snapshot or clone. Only a server outside the cluster should access the snapshot or clone.

Figure 54 shows a Windows server cluster in a SnapView configuration with a tape backup.




Figure 54 Example of a Windows cluster in a SnapView configuration with tape backup

Additional considerations

If MirrorView/A and/or SAN Copy sessions are running on the same source LUN as a SnapView snapshot session, these sessions share the same reserved LUN.

MirrorView/A and MirrorView/S limits

Table 58 MirrorView/A parameter limits for VNX series storage systems

Parameter ^a	Storage system limits		
	VNX5100	VNX5300, VNX5500	VNX5700, VNX7500
Primary images per mirror	1 required	1 required	1 required
Secondary images per mirror	1	1	1
Images (primary plus secondary) per storage system	128	256	256
Asynchronous consistency groups per storage system ^b	64 max	64 max	64 max
Mirrors per asynchronous consistency group ^c	32 max	32 max	64 max

- a. An image is either a LUN or a metaLUN. Since a metaLUN is a single entity, it counts as one image, even though it consists of multiple component LUNs.
- b. The consistency group limits are shared between MirrorView/A and MirrorView/S. The maximum number of MirrorView/A consistency groups plus MirrorView/S consistency groups per storage system is 64.
- c. The number of mirrors in asynchronous consistency groups is limited by the maximum number of mirrors for the storage system.

Note: Both the primary and secondary storage systems can each contain secondary mirror images, as well as primary mirror images (cross mirroring), provided they are for different mirrors. You cannot create a mirror of a SnapView clone or snapshot.

Table 59 MirrorView/S parameter limits for VNX series storage systems

Parameter ^a	Storage system limits			
	VNX5100, VNX5300	VNX5500	VNX5700	VNX7500
Primary images per mirror	1 required	1 required	1 required	1 required
Secondary images per mirror ^b	2	2	2	2
Images per storage system	128 max	256	512	1024
Synchronous consistency group per storage system ^c	64 max	64 max	64 max	64 max
Mirrors per synchronous consistency group ^d	32 max	32 max	64 max	64 max

- a. An image is either a LUN or a metaLUN. Since a metaLUN is a single entity, it counts as one image, even though it consists of multiple component LUNs.
- b. Only one secondary image for each mirror in a consistency group.
- c. The consistency group limits are shared between MirrorView/A and MirrorView/S. The maximum number of MirrorView/A consistency groups plus MirrorView/S consistency groups per storage system is 64.
- d. The number of mirrors in synchronous consistency groups is limited by the maximum number of mirrors for the storage system.

SAN Copy configuration limits

Table 60 SAN Copy parameter limits for VNX series systems

Parameter	Storage system limits		
	VNX5100, VNX5300	VNX5500	VNX5700, VNX7500
Maximum concurrent executing sessions (full and incremental) ^a per storage system	8 (4 per SP)	16 (8 per SP)	16 (8 per SP) (31) 32 (16 per SP) (32)
Maximum destination logical units ^b per session	50	50	100
Maximum incremental source logical units per storage system ^c	256	256	512
Maximum defined incremental sessions per source logical unit	8	8	8

- a. Concurrent SAN Copy sessions refers to SAN Copy sessions running at the same time.
- b. For a VNX for Block storage system, a logical unit is either a LUN or a metaLUN.
- c. For incremental SAN Copy, these limits include MirrorView/A images and SnapView snapshot source LUNs in addition to the incremental SAN Copy logical units. The maximum number of source LUNs assumes one reserved LUN assigned to each source LUN.

VNX RecoverPoint write Splitter usage

Supported storage systems

The following storage systems are supported with VNX RecoverPoint: VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500.

Supported connections

The RecoverPoint appliance connects to two ports per SP to achieve a HA configuration. Use any port *not* used for MirrorView/A or MirrorView/S. Host (server) connections to storage-system Fibre Channel or iSCSI data ports.

LUN support

A CX3 series storage system has a maximum of 1024 LUNs.

RecoverPoint mode support

Continuous data protection (CDP) and continuous remote replication (CRR) modes for the RecoverPoint target is supported.

[Table 61](#) lists the RecoverPoint source and target support for the MirrorView/A, MirrorView/S, SAN Copy, and SnapView applications.

Table 61 Interoperability with MirrorView/A, and MirrorView/S, SAN Copy, and SnapView

VNX for Block replication component	RecoverPoint source/target support	
	Source	Target (CDP & CRR)
MirrorView/A primary image	No ^a	No ^a
MirrorView/A secondary image	No	No
MirrorView/S primary image	No ^a	No ^a
MirrorView/S secondary image	No	No

Table 61 Interoperability with MirrorView/A, and MirrorView/S, SAN Copy, and SnapView

VNX for Block replication component	RecoverPoint source/target support	
	Source	Target (CDP & CRR)
SnapView clone source	Yes	Yes
SnapView clone	No	No
SnapView snapshot source	Yes	Yes
SnapView snapshot	No	No
Full SAN Copy source	Yes	Yes
Full SAN Copy destination	Yes	No
Incremental SAN Copy source	Yes	Yes
Incremental SAN Copy destination	Yes	No

- a. Although software does not currently prevent these configurations from being set up, these configurations are not currently supported.

Table 62 RecoverPoint limits

Parameter	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
NAS Consistency Groups per Cluster	NA	1	1	1	1
Max Source File Systems per Cluster	NA	2048	2048	2048	2048 14,336 (48)
Max source LUNs per Cluster	NA	NA	NA	NA	NA
Max Replicated Capacity per Cluster (TB)	600	600	600	600	600

CHAPTER 4

VNX R33 SW Configuration Rules

This chapter covers configuration rules for storage arrays running VNX R33. This software only runs on the following VNX models: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, VNX8000, VNX VG10, and VNX VG50. Topics include:

◆ Fibre Channel storage area network (SAN) configuration rules	190
◆ iSCSI network configuration rules.....	191
◆ FAST Cache for VNX storage systems	195
◆ Storage-system parameters for VNX series storage systems.....	196
◆ SnapView configuration rules	201
◆ VNX Snapshots	205
◆ MirrorView/A and MirrorView/S.....	206
◆ SAN Copy configuration limits	207
◆ SAN Copy configuration limits	207
◆ Multi-Core RAID (MCR).....	208
◆ Hot Spare Policy.....	208

Fibre Channel storage area network (SAN) configuration rules

Number of servers

Note: Highly-available (HA) servers contain an equal number of connections to each SP in the storage system.

For **VNX5200:**

Up to 1024 HA servers

For **VNX5400:**

Up to 1024 HA servers

For **VNX5600:**

Up to 1024 HA servers

For **VNX5800:**

Up to 2048 HA servers

For **VNX7600:**

Up to 4096 HA servers

For **VNX8000:**

Up to 8192 HA servers

The number of initiator records for the VNX series with VNX Operating Environment (OE) 05.33 for FC are shown in [Table 63](#)

Table 63 Initiator records for Fibre Channel

Storage system	Initiator records			
	Per FC port	Per FCoE port	Per SP	Per storage system
VNX5200	256	256	1024	2048
VNX5400	256	256	1024	2048
VNX5600	256	256	2048	2048
VNX5800	256	256	4096	4096
VNX7600	256	256	4096	8192
VNX8000	512	512	8192	16384

iSCSI network configuration rules

Note: Highly-available (HA) servers contain an equal number of connections to each SP in the storage system.

The number of initiator records for the VNX series with VNX OE for Block 05.33 for iSCSI are shown in [Table 64](#).

Table 64 Initiator records for iSCSI

Storage system	Initiator records	
	Per 1 Gb/s Ethernet iSCSI port	Per 10 Gb/s Ethernet iSCSI port
VNX5200	256	256
VNX5400	256	256
VNX5600	512	512
VNX5800	1024	1024
VNX7600	1024	1024
VNX8000	1024	1024

Each path used in a MirrorView or SAN Copy relationship between two storage system counts as an initiator for *both* storage systems.

Path rules for servers with failover software

Table 65 and Table 66 on page 194 list the number of paths supported from a server with failover software. For information on the storage systems supported by each type of failover software, see “Failover configuration rules” on page 140

Table 65 Number of paths supported from a server (page 1 of 3)

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
AIX	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	Native multipath failover ^c	Multipath	8	2
	DMP ^d	Multipath or Alternate Path ^e	8	16
HP-UX	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	Native multipath failover	Multipath	8	2
	PVLinks	Alternate Path	4	8
	DMP with VNX for Block driver	Multipath or Alternate Path ^e	8	16
Linux	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	Multipath failover (MPIO)	Multipath	4	8
	DMP with VNX for Block driver	Multipath or Alternate Path ^e	8	16
Net Ware	PowerPath	Multipath	N/A	16
	PowerPath SE ^b	Single Path	N/A	2
Open VMS	Native	Multipath	8	16

Table 65 Number of paths supported from a server (page 2 of 3)

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
Solaris	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	DMP with VNX for Block driver	Multipath or Alternate Path ^e	8	16
	Sun StorEdge Traffic Manager (SSTM)	Multipath	8	16
VMware ESX 3, ESXi 3, ESX Server 2.5,	Native	Alternate Path	8	16
Vmware ESX 4, ESXi 4	Native	Multipath	8	16
	PowerPath VE	Multipath	8	16

Table 65 Number of paths supported from a server (page 3 of 3)

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
Windows Server 2008, Windows Server 2008 R2	PowerPath	Multipath	8	16
	Multipath failover (MPIO)	Multipath	8	16

- Failover software that supports alternate path or multipath access to an SP also supports single-path access to an SP.
- For a description of PowerPath SE, see [“Failover configuration rules” on page 140](#).
- All multipath software is not supported for iSCSI configurations, and MirrorView/A, MirrorView/S, or SAN Copy configurations.
- See the VERITAS Volume Manager section of the E-Lab Interoperability Navigator for supported configurations.
- Multipath for DMP 4.1 or later; alternate path for DMP lower than 4.1.

Table 66 Number of paths supported from a server running Windows Server 2003 and Windows 2000

Server OS	Failover Software	Access type ^a	Number of paths supported from a server	
			To one SP	To one storage system
			VNX	
Windows Server 2003	PowerPath	Multipath	8	16
	PowerPath SE ^b	Single Path	1	2
	DMP with CLARiiON driver	Multipath or Alternate Path ^c	8	16
Windows 2000	PowerPath	Multipath	8	16
	PowerPath SE	Single Path	1	2
	DMP with CLARiiON driver	Multipath or Alternate Path ^c	8	16

- Failover software that supports alternate path or multipath access to an SP also supports single-path access to an SP.
- For a description of PowerPath SE, see [“Failover configuration rules” on page 140](#).
- Multipath for DMP 4.1 or later; alternate path for DMP lower than 4.1.

FAST Cache for VNX storage systems

Storage systems with VNX OE for Block 05.33 or later support an optional FAST Cache consisting of a storage pool of Flash disks that are configured to function as a FAST Cache. This cache provides low latency and high I/O performance without requiring a large number of Flash disks. It is also expandable while I/O to and from the storage system is occurring.

The FAST Cache is based on the locality of reference of the data set requested. A data set with high locality of reference that is frequently accessed is a good candidate for promotion to the FAST Cache. By promoting the data set to the FAST Cache, the storage system services any subsequent requests for this data faster from the Flash disks that make up the FAST Cache; thus, reducing the load on the disks in the LUNs that contain the data (the underlying disks). Applications such as file and OLTP (online transaction processing) have data sets that can benefit from the FAST Cache.

Note: The use of FAST cache (actually the enabler being installed) reduces the system’s DRAM cache.

The FAST Cache consists of one or more pairs of mirrored disks (RAID 1 type) and provides both read and write caching. For a read caching, the FAST Cache copies data off the disks being accessed into the FAST Cache. For protected-write applications, the FAST Cache effectively buffers the data waiting to be written to disk. In both cases, the workload is off-loaded from slow rotating disks to the faster Flash (SSD) disks in the FAST Cache. The performance boost provided by the FAST Cache varies with the workload and the cache size. The FAST Cache can be enabled for specific LUNs in RAID groups or all the LUNs in one or more pools. The FAST Cache should be disabled for Write Intent Log (WIL) LUNs or clone private LUNs. Enabling the FAST Cache for these LUNs is a misallocation of the FAST Cache and may reduce the effectiveness of the FAST Cache for other LUNs. [Table 67](#) lists the storage system type, Flash disk capacity, number of Flash disks in FAST Cache, and the maximum FAST Cache capacity.

Table 67 FAST Cache and disk configurations

Storage system	Number of Flash (SSD) disks in FAST Cache	Max FAST Cache Flash drive size, GB	Max FAST Cache/Secondary Cache, GB
VNX5200	6	200	600
VNX5400	10	200	1000
VNX5600	20	200	2000

Table 67 FAST Cache and disk configurations

VNX5800	30	200	3000
VNX7600	42	200	4200
VNX8000	42	200	4800

Storage-system parameters for VNX series storage systems

[Table 68](#) list the storage-system parameters for VNX series storage systems with VNX OE for Block 05.33 which include:

- ◆ Basic system parameters
- ◆ Memory and cache parameters
- ◆ Back end ports
- ◆ Data Mover I/O
- ◆ RAID group and RAID group LUN (traditional) parameters
- ◆ Pool and pool LUN parameters
- ◆ Storage group parameters
- ◆ Compression/Deduplication parameters
- ◆ Unisphere Quality of Service Manager parameters

Table 68 Storage system parameters (page 1 of 6)

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Basic system parameters						
Initiator records per FC port (max)	256	256	256	256	256	512
Initiator records per 1 Gb/s Ethernet iSCSI port (max)	256	256	256	512	1024	1024

Table 68 Storage system parameters (page 2 of 6)

Initiator records per 10 Gb/s Ethernet iSCSI port (max)	256	256	256	512	1024	1024
Initiator records per SP (max)	1024	1024	1024	2048	4096	8192
Front-end I/O modules/SP (max)	3	4	5	5	5	10
8 G/16G FC: Max I/O Ports/SP	12	16	20	20	20	36
1 GbE iSCSI: Max I/O Ports/SP	8	8	8	8	8	8
10 GbE iSCSI: Max I/O Ports/SP	8	8	8	8	8	8
FCoE: Max I/O Ports/SP	8	8	10	10	10	18
Disks (max)	125	250	500	750	1,000	1,500
Disks (min)	4	4	4	4	4	4
Memory and cache parameters						
	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Physical memory per SP, GB	16	16	24	32	64	128
MCC, GB	10.5	10.5	16.9	22.8	36.6	50.0
Max Write Cache with Data Services NOT Installed, GB	10.5	10.5	16.9	22.8	36.6	50.0

Table 68 Storage system parameters (page 3 of 6)

Max write cache size with Data Services installed, GB	4.5	4.5	9.7	13.7	31.9	50.0
Buffer cache, GB	4.0	4.0	4.7	6.0	14.3	31.8
FAST cache size per system (max), GB	600	1000	2000	3000	4200	4200
FAST cache disks (max)	6	10	20	30	42	42
Back End Ports						
6 G SAS BE Max ports/SP	2	2	6	6	6	16
6 G SAS BE Max SLICs/SP	NA	NA	1	1	1	4
Min SAS BE ports/SP	1	1	1	1	1	1
Max number of drives/BE port	125	250	250	250	250	250
Max number of enclosures/BE port	10	10	10	10	10	10
Max Block HA hosts (2 initiators/host)	1024	1024	1024	2048	4096	8192
	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Max FE ports/SP	12	16	20	20	20	36
Data Mover I/O						
Max number of data movers	3	4	4	6	8	8
Min number of data movers	1	1	1	1	2	2
Max FS + Snaps capacity per data mover	256	256	256	256	256	256

Table 68 Storage system parameters (page 4 of 6)

# of SLICs/DM	3	3	3	4	4	5
8 Gb FC Max ports/DM	4	4	4	4	4	4
# of FE SLICs/DM	2	2	2	3	3	4
1 GbE Max ports/DM	8	8	8	12	12	16
10 GbE Max ports/DM	4	4	4	6	6	8
RAID group and RAID group LUN (traditional) parameters						
Disks per private RAID group (max)	16	16	16	16	16	16
RAID groups (max)	62	62	125	250	500	1000
LUNs (max)	2048	2048	2048	4096	4096	8192
metaLUNs (max)	512	512	512	512	1024	2048
metaLUN volumes (max) ^a	512	512	512	512	1024	2048
LUNs per metaLUN component (max)	32	32	32	32	32	32
Components per metaLUN (max)	8	8	8	8	16	16
LUNs in a MetaLUN (max)	256	256	256	256	512	512
Max total consumables	2048	2048	2048	4096	4096	8192
RAID types ^b	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk	6, 5, 3, 1/0, 1, individual disk

Table 68 Storage system parameters (page 5 of 6)

Pool and pool LUN parameters						
Max pool capacity, TB	424	861	1,736	2,611	3,486	3,486
Pool LUNs per storage system (max)	1,000	1,000	1,000	2,000	3,000	4,000
Pool LUN size (max)	256 TB	256 TB	256 TB	256 TB	256 TB	256 TB
Disks in all pools (max)	121	246	496	746	996	996
Disks per pool (max)	121	246	496	746	996	996
Pools (max)	15	15	20	40	40	60
LUNs per pool (max)	1,000	1,000	1,000	2,000	3,000	4,000
RAID types	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0	6, 5, 1/0
Max number of disks that can be added to a pool at a time	80	80	120	120	120	180
Storage group parameters						
Storage groups (max)	256	256	256	512	1024	2048
LUNs per storage group (max)	4096	4096	4096	4096	4096	4096
Compression/Deduplication parameters						
Concurrent compression operations	20	20	20	20	32	40
Concurrent deduplication migrations	8	8	8	8	8	8

Table 68 Storage system parameters (page 6 of 6)

Fixed Block de-dupe sessions	4	4	4	4	4	4
Max number of compressed LUNs involving migration per storage system	16	16	16	16	24	24
Max number of compressed LUNs	1000	1000	1000	2000	3000	4000
Unisphere Quality of Service Manager parameters						
Policies per storage system (max)	10	10	10	10	10	10
I/O classes per storage system (max)	64	64	64	64	64	64
LUNs per I/O class (max)	128	128	128	128	128	128
I/O classes per cruise control policy (max)	2	2	2	2	2	2
I/O classes per limit or queue depth policy (max)	32	32	32	32	32	32
Scheduled events (max)	10	10	10	10	10	10

- a. Max Pool LUN is 16 TB. Max RG LUN can be larger than 16 TB. Larger LUNs can be configured provided host operating system can consume those LUNs.
- b. For best practices, use RAID 6 for less efficient drives and RAID 5 or RAID 1/0 for more efficient drives.

SnapView configuration rules

Note: *LUN* refers to *LUN* or *metaLUN* (not a component of a metaLUN).

IMPORTANT

Do not allocate SnapView reserved LUNs or clones from the same RAID Group as the Source LUNs that will be replicated. Doing so impacts performance severely.

Supported storage systems

The following storage systems running R33 are supported for SnapView: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, and VNX8000.

Source LUNs for clones

Any source LUN or thin LUN can be cloned, *except for*:

- ◆ Hot spare LUNs
- ◆ Clone LUNs (LUNs participating in any clone group as either a source LUN or a clone LUN)
- ◆ Snapshot LUNs
- ◆ Private LUNs (LUNs reserved as clone private LUNs in a reserved LUN pool or in a write intent log)
- ◆ A LUN in the process of conversion

Table 69 on page 202 lists the SnapView clone parameters and their limits for VNX storage systems.

Table 69 SnapView supported limits for VNX series storage systems

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Clones (clone images) ^a						
Clones per storage system	256	256	1024	2048	2048	2048
Max clone source LUNs	128	128	512	1024	1024	1024
Clones per consistent fracture	32	32	32	32	64	64

Table 69 SnapView supported limits for VNX series storage systems (continued)

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Max clones in a clone group	8	8	8	8	8	8
Clone private LUNs^{b, c}						
Per storage system (required)	2	2	2	2	2	2
Clone source LUNs						
Per storage system	64	64	128	256	512	1024

a. The source LUN and MirrorView/S primary and secondary images no longer count towards the clone image limit.

b. A thin LUN cannot be a clone private LUN.

c. The minimum size for a clone private LUN is 1 GB.

Table 70 lists the SnapView snapshot parameters and their limits for VNX storage systems.

Table 70 Maximum SnapView snapshot limits for VNX series storage systems

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Snapshots^a						
Per source LUN	8	8	8	8	8	8
SnapView sessions^a						
Per source LUN	8	8	8	8	8	8
Per storage system	512	512	512	512	1024	2048
Source LUNs in consistent session	32	32	32	32	64	64
Source LUNs						
Max Snapshot Source LUNs	256	256	256	256	512	512
Per storage system	512	512	512	512	1024	2048

a. The limits for snapshots and sessions include SnapView snapshots or SnapView sessions, as well as reserved snapshot or reserved sessions used in other applications, such as and MirrorView/A and SAN Copy (incremental sessions).

Note: You can create a snapshot of a clone, but you cannot create a clone of a snapshot. You can create a snapshot of a mirrored LUN, but you cannot create a clone of a mirrored LUN.

Supported operating systems

All supported operating systems work with SnapView. The server that uses the source LUN and the server that uses the replica (snapshot or clone) must run the same operating system.

Number of servers

Two or more, where one or more servers have access to the source LUN and one or more servers have access to the replica (snapshot or clone). EMC supports placing a clone or snapshot in the same storage group as its source LUN *only for a server running*

- ◆ Windows Server 2008, Windows Server 2008 R2, Windows Server 2008 Hyper-V, Windows Server 2003, or Windows 2000.
- ◆ VMware ESX 4, VMware ESXi 4, VMware ESX 3, VMware ESXi 3, VMware ESX Server 2.5 with the clone or snapshot is presented to a different virtual machine than its source LUN.
- ◆ Other operating systems, if you use RM/Local, RM/SE, or PowerPath Volume Manager to put the clone or snapshot in the storage group.

These operating systems and software products provide “same server access” to the clone, snapshot, and the source LUN. For information on using these operating systems or software products, refer to the documentation for the operating system or product.

VNX Snapshots

EMC® VNX Snapshots™ is a storage system-based software application that allows you to create snapshots of pool-based LUNs. A snapshot is a virtual point-in-time copy of a LUN and takes only seconds to create. You can use the snapshot for data backups, decision support, data modeling, software application testing, or as a base for temporary operations on the production data without damaging the original data on the source LUN.

Table 71: VNX Snapshot parameters

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Snapshots						
Per storage system	8,000	8,000	8,000	16,000	24,000	32,000
Per primary LUN	256	256	256	256	256	256
Per consistency group	64	64	64	64	64	64
VNX Snapshot Source LUNs						
Per System	1,000	1,000	1,000	2,000	3,000	4,000
Per Consistency group	64	64	64	64	64	64

Table 71: VNX Snapshot parameters

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Consistency groups						
Per storage system	128	128	128	128	256	256
Snapshot Mount Points						
Per storage system	1000	1000	1000	2000	3000	4000
Concurrent restore operations						
Per storage system	128	128	128	256	512	512

MirrorView/A and MirrorView/S

Table 72: MirrorView/A and MirrorView/S parameters

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
MirrorView/A						
Max Mirrorview/A objects	256	256	256	256	256	256
MirrorView/A consistency groups	64	64	64	64	64	64
MirrorView/A mirrors per consistency group	32	32	32	32	64	64
MirrorView/S						

Table 72: MirrorView/A and MirrorView/S parameters (continued)

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Max MirrorView/S objects	128	128	128	256	512	1024
MirrorView/S consistency groups	64	64	64	64	64	64
MirrorView/S mirrors per consistency group	32	32	32	32	64	64

SAN Copy configuration limits

Table 73 SAN Copy limits

Parameter	VNX5200	VNX5400	VNX5600	VNX5800	VNX7600	VNX8000
Destinations per Source	50	50	50	50	100	100
Max Concurrent SAN Copy Sessions per array	8	8	8	16	32	32
Max Incremental SAN Copy sources	256	256	256	256	512	512
Incremental SAN Copy sessions per Source	8	8	8	8	8	8

Multi-Core Cache (MCC)

In R33 Multi-Core Cache replaces the old FLARE SP cache. The internal caching algorithm has been rewritten for multi-core processors. As a result, MCC is completely multi-core scalable. In addition, the number of tunable parameters is reduced. Cache page size is now locked at 8KB. Cache can be enabled on the array or on a single LUN.

As of R33:

- ◆ SP cache is dynamically allocated from the system, instead of manually by a user

- ◆ Read caching is always enabled on arrays and LUNs
- ◆ SP/LUN cache user controls are removed, except for the write cache enable/disable options
- ◆ Cache statistics options are changed

Multi-Core RAID (MCR)

In R33, MCR replaces the FLARE RAID engine. MCR is multi-core scalable and active/active backend architecture. MCR includes the following:

- ◆ Hot Spare Policy
- ◆ Portable Drives
- ◆ Permanent Drive Sparing
- ◆ Symmetric Active/Active

Hot Spare Policy

The VNX series now supports a new hot spare policy where any unbound disk is available for use as a hot spare. The hot spare policy defaults to the EMC recommended best practices for disk type and capacity but you can customize the policy to either increase or decrease the ratio of disks that you want to keep unused as hot spares.

The hot spare policy is enforced whenever disks are provisioned, such as when creating RAID Groups, pools, or FAST Cache. If a policy is violated, alerts automatically display in Unisphere or the Command Line Interface (CLI) to identify the disks.

You also have options of:

- Using a copy-to operation to move data back to the original location after a replacement drive is inserted into the failed slot.
- Physically moving the permanent spare disk to the failed location.

Portable Drives

Another new feature of MCR is the ability to move drives freely around the array. Any drive can be moved to any slot (except for the vault) and remain part of a RAID group or a pool that it was in. Entire RAID groups can be moved.

Permanent Drive Sparing

When a drive fails, the system uses a policy to find a replacement drive. MCR introduces the notion of permanent spare and eliminates the need for the user to designate hot spares. The system will simply use any spare drive as a possible replacement drive.

Symmetric Active/Active

Although VNX arrays have been Active/Active from the beginning--that is, any LUN could be read and written to by either SP at any time--in the SCSI ALUA protocol, paths contain indicators of LUN ownership, so that the host could prefer one SCSI path over another. This was used as a remedy to the fact that non-owning SP took more overhead to access a LUN belonging to the other SP. With the new feature, that extra overhead has been eliminated, regardless of which SP accessed the LUN, and all paths indicate that they have been Optimized. While this feature improved all LUNs' performance on the previously non-optimized path, the overhead was only eliminated for RAID group LUNs, and therefore only RAID group LUNs may report that all paths are Optimized.

Deduplication and compression

Block Deduplication

VNX2 systems introduced VNX Block Deduplication. In general, deduplication is the process of identifying duplicate data contained within a set of block storage objects and consolidating it such that only one actual copy of the data is used by many sources. This feature can result in significant space savings, depending on the nature

of the data. VNX Block Deduplication utilizes the fixed block deduplication method with a set size of 8 KB to remove redundant data from a dataset. Block Deduplication is run post-process on the selected dataset.

Block Deduplication is enabled on a per pool LUN basis. Once deduplication is enabled on an existing pool LUN, the LUN must be migrated using a background process into the deduplication container. The data must reside in the container before the deduplication process can run on the dataset. During this migration the deduplication state for the LUN will show as Enabling and the migration process can be viewed on the Deduplication tab within the LUN Properties window. Once the migration is complete, the LUN will be Thin and deduplication is enabled. To enable deduplication on a Classic LUN, first create a deduplication-enabled LUN within a pool and use LUN migration to migrate the data to it.

For applications requiring consistent and predictable performance, EMC recommends utilizing Thick pool LUNs or Classic LUNs. If Thin LUN performance is not acceptable, then block deduplication should not be used.

VNX Block deduplication runs in the background on the contents of a deduplication container post process, 12 hours after the previous deduplication run completed on the pool. Each pool in the system has an independent deduplication container and process unique to the pool. Only the contents in a deduplication container are compared against each other.

When the 12-hour timer expires and the deduplication process is set to run, the number of running deduplication processes on the current SP is checked. A maximum of 3 deduplication processes per SP may run at a time, and if an open slot is not found, the deduplication process is queued until one becomes available. The SP the process runs on is based on the SP that the first deduplicated LUN was bound on.

While the basic purpose of Block Deduplication is to eliminate redundant data and to save storage space, EMC suggests using deduplication with multicore FAST Cache and FAST VP to achieve better efficiency of these features.

IMPORTANT

VNX Block Deduplication is not compatible with compression because the current implementation of compression is not compatible with the Deduplication container.

VNX Block Compression

VNX series storage systems with VNX OE for Block 05.33 and both the Compression and Thin Provisioning enablers provide the option of compressing the data on the LUN to free up storage space. Compression analyzes the data on a disk and applies algorithms that reduce the size of repetitive sequences of bits that are inherent in some types of files. Unlike de-duplication, compression does not depend on multiple copies of the same data. The storage system performs compression operations in the background while continuing to service the host I/O.

In general, compression is a process which attempts to reduce the total space used by a dataset. VNX Block Compression works in 64 KB increments to reduce the storage footprint of a dataset and provide savings to the user.

VNX Block Compression can be enabled on any LUN type on the system. Once compression is enabled on a pool-based LUN, and initial compression process is run and the LUN becomes Thin. If Thin LUN performance is not acceptable, then Block Compression should not be used.

VNX Block Compression processes data in 64 KB increments. It will only compress data if at least 8 KB of the 64 KB can be saved. If the resulting savings from compression is less than 8 KB within the 64 KB chunk, the data is written to the LUN uncompressed.

Compression of a RAID group or thick LUN converts the LUN to a thin LUN. The compressed LUN remains a thin LUN even after you decompress it, but the LUN's host-visible size is equal to the size that you specified when you created the LUN. You cannot compress a newly created thick LUN without data on it. For these reasons, the recommended procedure for creating a compressed LUN is to create the LUN as a thin LUN, and once created, turn on compression for the LUN.

For additional resources on VNX2 deduplication and compression, see:

- ◆ *EMC VNX Deduplication and Compression: VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, & VNX8000 Maximizing effective capacity utilization* [VNX2 white paper]
- ◆ *EMC VNX Block Deduplication: Basic Overview* [EMC Community Network document]

CHAPTER 5

Rail Kits for VNX Family Systems

This chapter covers the field-install rail kits for VNX systems. Topics include:

- ◆ Rail kits..... 214

Rail kits

EMC sells rail kits for mounting VNX storage-system enclosures in 19-inch NEMA cabinets/racks and Telco racks. Standard NMEA racks are listed in [Table 74](#) and [Table 75](#). Telco racks are listed in [Table 77 on page 215](#) and [Table 78 on page 215](#).

Standard NEMA racks

Table 74 VNX5100 and VNX5300 rail kits

Model Number	For
VNX-DPE15ADJS	Adjustable rail kit for 3U 15 drive DPE
VNX-DPE25ADJS	Adjustable rail kit for 3U 25 drive DPE
VNX-DAE15ADJS	Adjustable rail kit for 3U 15 drive DAE
VNX-DAE25ADJS	Adjustable rail kit for 2U 25 drive DAE
VNX-DMEADJS	Adjustable rail kit for 2U Data Mover Enclosure
VNX-SPS-ADJS	Adjustable rail kit for 1U SPS
VNX-CS-ADJS	Adjustable rail kit for 1U Control Station

Table 75 VNX5500, VNX5700, and VNX7500 rail kits

Model Number	For
VNX-DPE15ADJ	Adjustable rail kit for 3U 15 drive DPE
VNX-DPE25ADJ	Adjustable rail kit for 3U 25 drive DPE
VNX-DAE15ADJ	Adjustable rail kit for 3U 15 drive DAE
VNX-DAE25ADJ	Adjustable rail kit for 2U 25 drive DAE
VNX-DMEADJ	Adjustable rail kit for 2U Data Mover Enclosure
VNX-SPS-ADJ	Adjustable rail kit for 1U SPS
VNX-SPE2UADJ	Adjustable rail kit for 2U SPE
VNX-CS-ADJ	Adjustable rail kit for 1U Control Station

Table 76 VNX5200, VNX5400, VNX5600, VNX5800, VNX7600 and VNX8000 rail kits

Model Number	For
VNXBCS1UADJ	Adjustable VNXB 1U CS rail kit
VNXBDAE2UADJ	Adjustable VNXB 2U DAE rail kit
VNXBDME2UADJ	Adjustable VNXB 2U DME rail kit
VNXBDAE3UADJ	Adjustable VNXB 3U DAE rail kit
VNXBDPE3UADJ	Adjustable VNXB 3U DPE rail kit
VNXBSPE4UADJ	Adjustable VNXB 4U SPE rail kit

TELCO racks

Table 77 VNX5100 and VNX5300 TELCO rail kits

Model Number	For
DPE-TELCO-3US	DPE Telco Rail Kit for the 3U DPE
FRAME-T25DAES	Telco Rail kit for the 2U DAE
FRAME-T-PDAES	Telco Rail kit for the 3U DAE
SPST-2GBS	Telco Rail Kit for the 1U SPS
SPE-TELCO-2US	Telco Rail Kit for the 2U AC DME
FRAME-T-CSS	Telco Rail Kit for the Control Station

Table 78 VNX5500, VNX5700, and VNX7500 TELCO rail kits

Model Number	For
DPE-TELCO-3U	Telco Rail Kit for the 3U DPE
FRAME-T-PDAEV	Telco Rail Kit for the 3U DAE
FRAME-T25DAE	Telco Rail Kit for the 2U 25 2.5" Drive DAE

Table 78 VNX5500, VNX5700, and VNX7500 TELCO rail kits

Model Number	For
SPE-TELCO-2UV	Telco Rail Kit for the 2U AC SPE and DME
SPST-2GBV	Telco Rail Kit for the 1U SPS
FRAME-T-CS	Telco Rail Kit for the Control Station

Table 79 VNX5200, VNX5400, VNX5600, VNX5800, and VNX7600 TELCO rail kits

Model Number	For
VBTELCO1UCS	Telco Rail Kit for the 1U CS
VBTELCO2UDME	Telco Rail Kit for the 2U DME
VBTELCO3UDPE	Telco Rail Kit for the 3U DPE
VBTELCO3UDAE	Telco Rail Kit for the 3U DAE
VBTELCO2UDAE	Telco Rail Kit for the 2U DAE

CHAPTER 6

VNX and Fibre Channel Switches

This chapter covers VNX cabling with and without switches. Topics include:

- ◆ Fibre Channel switch zoning..... 218

Fibre Channel switch zoning

Zoning spans the entire fabric. EMC requires single-initiator zoning using the World Wide Port Name (WWPN) of the HBA port and the WWPNs of the SP ports. One server can have only one path to each SP, unless the server is running software that supports multipath or alternate path access to the storage-system SP (see [“Paths from a server to an SP” on page 113](#)). For information on configuring zones, refer to the switch manuals and release notes.

IMPORTANT

Be sure to document all zoning parameters, including the WWPNs of HBA ports and SP ports, for future reference.

Remote mirror zoning

The [“Fibre Channel storage area network \(SAN\) configuration rules” on page 95](#) discusses HBA or CNA port zoning. If an HBA is using one or more SP ports for normal I/O and another SP port (VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500) for normal I/O and remote mirror data, use one switch zone for the HBA port and the normal I/O ports. Use another switch zone for the HBA port and the remote mirror data ports.

Note: For VNX MirrorView port information, see [“VNX series MirrorView ports” on page 94](#).


QuickLoop

QuickLoop emulates an Arbitrated Loop in a Fabric.

- QuickLoop is supported only on DS-8B, DS-16B, and DS-16B2 (1-Gb mode) switches.
- One QuickLoop can be present on each switch or span two switches.
- Servers and storage systems within a QuickLoop must be configured for an Arbitrated Loop.

Zoning examples

Figure 55 shows Single-Initiator zoning using multiple HBA-port servers, multiple FC switches, and multiple storage systems.



Legend (A) = SP A and (B) = SP B

A component failure in the solid (red) line path prevents I/O from any server. PowerPath or DMP on that server allows all access via the dashed (blue) line path.

Figure 55 Single-Initiator zoning with multiple HBA-port servers

Storage group assumptions

- ◆ Server A has a storage group on storage systems W and X.
- ◆ Server B has a storage group on storage systems V and X.
- ◆ Server C has a storage group on storage systems V, W, and X.

- ◆ Servers D and E have a storage group on storage systems Y and Z.

MirrorView assumptions

- ◆ Storage systems W and Y are running MirrorView/S or MirrorView/A.

SP ports used for host I/O

- ◆ Server A uses all connected SP ports on storage systems W and X for host I/O.
- ◆ Server B uses all connected SP ports on storage systems V and X for host I/O.
- ◆ Server C uses SP port 0 on storage systems V, W, X for host I/O.
- ◆ Server D uses all connected SP ports on storage system Y and SP port 0 on storage system Z for host I/O.
- ◆ Server E uses all connected SP ports on storage systems Y and Z for host I/O.

[Figure 56 on page 221](#) shows Single-Initiator zoning using dual HBA-port servers with FC fabric connected to a single system.




Figure 56 Single-initiator zoning with dual HBA-port servers

Table 80 lists the storage processor, switch zone, storage processors and ports in the zone, and storage processor mirror ports on storage system connections.

Table 80 Zone contents for storage-system Mirror Ports

SP	Switch zone	Participants in zone by WWPN	Connections
SP A	1	SP 2 port 1, SP 6 port 3	SP A mirror ports on storage systems W and Y
SP B	2	SP 3 port 1, SP7 port 3	SP B mirror ports on storage systems W and Y

Table 81 lists the server, switch zone, host bus adapter, storage processor, and ports in the zone, and host bus adapter and storage processor mirror and non-mirror port connections.

Table 81 Zone contents for server HBAs (page 1 of 2)

Server	Switch zone	Participants in zone by WWPN	Connections
Server A	3	HBA 0, SP 2 port 0, SP 4 port 0, SP 5 port 1	HBA 0 and SP non-mirror ports
	4	HBA 0, SP 3 port 1	HBA 0 and SP mirror port
	5	HBA 1, SP 3 port 0, SP 4 port 1, SP 5 port 0	HBA 1 and SP non-mirror ports
	6	HBA 1, SP 2 port 1	HBA 1 and SP mirror port
Server B	7	HBA 2, SP 0 port 0, SP 1 port 1, SP 4 port 0, SP 5 port 1	HBA 2 and SP non-mirror ports
	8	HBA 3, SP 0 port 1, SP 1 port 0, SP 4 port 1, SP 5 port 0	HBA 3 and SP non-mirror ports

Table 81 Zone contents for server HBAs (page 2 of 2)

Server	Switch zone	Participants in zone by WWPN	Connections
Server C	9	HBA4, SP0 port 0, SP2 port 0, SP4 port 0	HBA4 and SP non-mirror ports
	10	HBA4, SP3 port 1	HBA4 and SP mirror port
	11	HBA5, SP1 port 0, SP3 port 0, SP5 port 0	HBA5 and SP non-mirror ports
	12	HBA5, SP2 port 1	HBA5 and SP mirror port
Server D	13	HBA6, SP6 port 0, SP6 port 2, SP7 port 1, SP7 port 3, SP8 port 0	HBA6 and SP non-mirror ports
	14	HBA6, SP7 port 3	HBA6 and SP mirror port
	15	HBA7, SP6 port 0, SP6 port 2, SP7 port 1, SP7 port 3, SP8 port 0	HBA7 and SP non-mirror ports
	16	HBA7, SP7 port 3	HBA7 and SP mirror port
	17	HBA8, SP6 port 1, SP6 port 3, SP7 port 0, SP7 port 2, SP9 port 0	HBA8 and SP non-mirror ports
	18	HBA8, SP6 port 3	HBA8 and SP mirror port
	19	HBA9, SP6 port 1, SP6 port 3, SP7 port 0, SP7 port 2, SP9 port 0	HBA9 and SP non-mirror ports
	20	HBA9, SP6 port 3	HBA9 and SP mirror port
Server E	21	HBA10, SP6 port 0, SP6 port 2, SP7 port 1, SP7 port 3, SP8 port 0, SP9 port 1	HBA10 and SP non-mirror ports
	22	HBA10, SP7 port 3	HBA10 and SP mirror port
	23	HBA11, SP6 port 1, SP6 port 3, SP7 port 0, SP7 port 2, SP8 port 1, SP9 port 0	HBA11 and SP non-mirror ports
	24	HBA11, SP6 port 3,	HBA11 and SP mirror port

Figure 57 shows Single-Initiator zoning using dual and single HBA-port servers with FC fabric connected to a single system.




Figure 57 Single-initiator zoning with dual HBA-port and single HBA-port servers connected to a VNX series system

Configuration goals

- Connect both ports of SPs so different servers can use the SP front end ports.
- Cross port 1 on each SP to the other fabric so the configuration is ready for the use of multipath PowerPath, or alternate-path DMP in the future.
- Use switch zoning to ensure that each server can see only one path to an SP.

Zone contents [Table 82](#) lists the switch zone, host bus adapter and storage processor adapter in the zone, and the storage group that the server has in the system

Table 82 Zone contents

Switch zone number	Participants in zone by WWPN	Server
1	HBA 0, SP A0	Server A has a storage group on the VNX system.
2	HBA 1, SP B0	
3	HBA 2, SP B1	Server B has a storage group on the VNX system.
4	HBA 3, SP A1	
5	HBA 4, SPA1, SP B0	Server C has a storage group on the VNX system.

Fibre Channel and FCoE switch topology rules

For information on EMC switch topology and specific valid configurations, refer to the following documents:

- *EMC Networked Storage Topology Guide*
- *E-Lab Interoperability Navigator*

CHAPTER 7

VNX Cable Types

This chapter covers the optical and Twin Ax cables used with the VNX systems. Topics include:

- ◆ VNX cable types 228

VNX cable types

Optical cables

These models consist of a 50/125 micron, 2mm multi-mode optical cable. Depending on the length, the cable assembly consists of either riser or plenum type cable as listed in [Table 83](#). Both ends of the cable will have LC style connectors, which comply with the IEC 61754-20 connector family standard.

Table 83 Optical OM3 cables for 8 Gb, 4 Gb, or 2 Gb operation^a

Cable Length	LC-LC Multi-mode OM3 fibre optic cables		
	VNX5100/5300	VNX5500/VNX5700/ VNX7500	VNX5200/VNX5400/ VNX5600/VNX5800/ VNX7600/VNX8000
1 meter	VNX-OM3-1MS	VNX-OM3-1M	VNXB-OM3-1M
3 meter	VNX-OM3-3MS	VNX-OM3-3M	VNXB-OM3-3M
5 meter	VNX-OM3-5MS	VNX-OM3-5M	VNXB-OM3-5M
10 meter	VNX-OM3-10MS	VNX-OM3-10M	VNXB-OM3-10M
30 meter	VNX-OM3-30MS	VNX-OM3-30M	VNXB-OM3-30M
50 meter	VNX-OM3-50MS	VNX-OM3-50M	VNXB-OM3-50M
100 meter	VNX-OM3-100MS	VNX-OM3-100M	VNXB-OM3-100M

a. The VNX5100/5300 & VNX5500/5700/7500 cables are the same physical cable. Different models are used for discount classes and pricing purposes.

TwinAx cables

These models consist of a shielded, quad construction style cable with a 100 Ohm differential as listed in Table 84. Both ends of the cable have SFP+ style connectors that comply with SFF-8431 and SFF-8472 standards. The transmit and receive ends of the cable have active components to facilitate the transmission of 8 Gigabit fibre channel and/or 10 Gigabit Ethernet protocols. The use of DC blocking capacitors on the receiver is required per the SFF-8431 standard

Table 84 Active TwinAX cable for FCoE or 10Gb iSCSI operation^a

Cable Length	Description	SFP+ to SFP+ Active 10Gb/8Gb TwinAx cables		
		VNX5300	VNX5500 VNX5700 VNX7500	VNX5200, VNX5400 VNX5600, VNX5800 VNX7600, VNX8000
1 meter	SFP+ TO SFP+ 1M ACTIVE 8G/10G CABLE	SFP-TWNACT-1MS	SFP-TWNACT-1M	VBSFP-TWAX-1M
3 meters	SFP+ TO SFP+ 3M ACTIVE 8G/10G CABLE	SFP-TWNACT-3MS	SFP-TWNACT-3M	VBSFP-TWAX-3M
5 meters	SFP+ TO SFP+ 5M ACTIVE 8G/10G CABLE	SFP-TWNACT-5MS	SFP-TWNACT-5M	VBSFP-TWAX-5M

a. The VNX5300 & VNX5500/5700/7500 cables are same physical cable. Different models are used for discount classes and pricing purposes.The VNX5100 system does not support FCoE TwinAx technology.

IMPORTANT

Customers with TwinAX I/O modules must use supported Active TwinAX cables.

SAS cables

DPE to DAE cables

Table 85 VNX5100, VNX5300, VNX5500 DPE to DAE SAS cables

Cable Length	
	Model
Mini-SAS to Mini-SAS 2M cable	MSAS-MSAS-2M
Mini-SAS to Mini-SAS 5M cable	MSAS-MSAS-5M

Table 86 DPE/SPE to DAE SAS cables

Description	VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, VNX8000 Model Number
ONE PR 2M MINI SAS HD TO MINI SAS CBLs	VBMSASHDMSAS2
ONE PR 5M MINI SAS HD TO MINI SAS CBLs	VBMSASHDMSAS5
ONE PR 8M MINI SAS HD TO MINI SAS CBLs	VBMSASHDMSAS8

SPE to DAE cables

Table 87 SPE to DAE SAS cables

Cable Length		
	VNX5700 VNX7500	VNX5200, VNX5400 VNX5600, VNX5800 VNX7600, VNX8000
Mini-SAS HD to Mini-SAS 2M cable	MSASHD-MSAS2	VBMSASHDMSAS2
Mini-SAS HD to Mini-SAS 5M cable	MSASHD-MSAS5	VBMSASHDMSAS5
Mini-SAS HD to Mini-SAS 8M cable	MSASHD-MSAS8	VBMSASHDMSAS8

DAE to DAE cables

Table 88 DAE to DAE SAS cables for VNX1 series

SAS Cable Length		
	VNX5100, VNX5300, VNX5500, VNX5700 VNX7500	Notes
Mini-SAS to Mini-SAS 2M cable	MSAS-MSAS-2M	
Mini-SAS to Mini-SAS 5M cable	MSAS-MSAS-5M	
Mini-SAS to Mini-SAS 10M cable	MSAS-MSAS-10	
Mini-SAS to Mini-SAS 2M hi-flex cables	SAS-FLEX-2M	For use with DAE7S only
Mini-SAS to Mini-SAS 3M hi-flex cables	SAS-FLEX-3M	For use with DAE7S only
Mini-SAS to Mini-SAS 8M hi-flex cables	SAS-FLEX-8M	For use with DAE7S only

Table 89 DAE to DAE SAS cables for VNX2 series

Description	VNX5200, VNX5400, VNX5600, VNX5800, VNX7600, VNX8000
ONE PR OF 2M MINI SAS TO MINI SAS CBLS	VBMSAS-MSAS2M
ONE PR OF 5M MINI SAS TO MINI SAS CBLS	VBMSAS-MSAS5M
ONE PR OF 10M MINI SAS TO MINI SAS CBLS	VBMSAS-MSAS10
HI-FLEX 2M MINISAS to MINISAS CBLS	VBSAS-FLEX2M
HI-FLEX 3M MINISAS to MINISAS CBLS	VBSAS-FLEX3M
HI-FLEX 5M MINISAS to MINISAS CBLS	VBSAS-FLEX5M
HI-FLEX 8M MINISAS to MINISAS CBLS	VBSAS-FLEX8M

INDEX

A

- Add blade 89
- Add blade enclosure 89
- Add Control Station 89
- Add DAE 89
- Add disk 89
- Add I/O module to blade 89
- Add I/O module to SP 89
- Add SPS 89
- Asynchronous login permitted 131

B

- Back End Ports 198
- Back-end data ports per SP 171
- Basic system parameters 196
- Block Compression 211
- Block Configuration 12
- Block Deduplication 209
- Block to Unified upgrades 90
- Boot LUN type 132
- Booting from SAN 131
 - Boot LUN isolation 137
 - Cluster software 138
 - Combination of local and remote pagefile locations 135
 - Failover mode 133
 - General considerations 132
 - General notes on pagefiles 135
 - Local pagefile 133
 - Maximum boot servers 132
 - Memory dump file support (Windows) 134
 - operating system specific boot information 135
 - pagefiles location considerations 133
 - Remote pagefile 134
 - Replication and boot LUNs 136
 - Server boot LUNs 133
 - Support for boot LUNs with clusters 137

C

- Clustered server configuration rules 176

- Components per metaLUN (max) 172, 199
- Compression for systems with VNX OE 05.31 and 05.32 169
- Compression for VNX storage systems 211
- Compression/Deduplication parameters 200
- Concurrent compression operations 174, 200
- Control Station (CS) types 26

D

- DAE Configuration Rules 24
- DAE to DAE cables 231
- DAE to DAE SAS cables 231
- DAE5S 21
- DAE6S 22
- DAE7S 22
- Data in place conversions 91
- Data Mover I/O 198
- Data Mover types 27
- DC power 14
- Dense storage 14
- Disk array enclosure (DAE) types 21
- Disk array enclosure with 3.5 drives and 40U deep rack 23
- disk power savings 128
- Disks (max) 171, 197
- Disks (min) 171, 197
- Disks in all pools (max) 173, 200
- Disks per pool (max) 173, 200
- Disks per RAID group (max) 172, 199
- domain security settings 131
- DPE 15
- DPE to DAE cables 230
- DPE to DAE SAS cables 230
- DWDM 95

F

- Fabric 95
- Failover configuration rules 140
- FAST Cache 168, 195
- FAST Cache capacity 168, 195

FAST cache disks (max) 198
FAST cache size per SP 198
FCoE gateway configuration guidelines 99
FCoE I/O module 100
FCoE support for gateways 98
FE logical port identifiers 30
Fibre Channel 101
Fibre Channel and FCoE switch topology rules 90, 155,
157, 160, 184, 211, 225
Fibre Channel over Ethernet 101
Fibre Channel Storage Area Network (SAN)
Configuration Rules 95, 164, 190
File Configuration 12
Front-end data ports per SP 171, 197

H

Hardware upgrades 89
HBA 95
host connection 95, 164, 190
Hot Spare Policy 208

I

I/O classes per cruise control policy (max) 174, 201
I/O classes per limit or queue depth policy (max) 174,
201
I/O classes per storage system (max) 174, 201
I/O module types 29
Initiator records 164, 165, 190, 191
Initiator records per 1 Gb/s Ethernet iSCSI port 171,
196
Initiator records per 10 Gb/s Ethernet iSCSI port 171,
197
Initiator records per 10GBase-T iSCSI/IP port 171
Initiator records per SP 171, 197
Initiator records per storage system 171
iSCSI network configuration rules 110
ISL 96

L

local users support 130
LUN or metaLUN size (max) 172, 199
LUNs (max) 172, 199
LUNs per I/O class (max) 174, 201
LUNs per metaLUN component (max) 172, 199
LUNs per pool (max) 173, 200
LUNs per RAID group (max) 172

LUNs per storage group (max) 173, 200

M

Maximum concurrent executing sessions (full and
incremental) per storage system 185
Maximum destination logical units per session 185
Maximum incremental source logical units per storage
system 185
Maximum number of reserved LUNs 147
Memory and cache parameters 197
metaLUNs 175
metaLUNs (max) 172, 199
MirrorView Fibre Channel FE ports 94
MirrorView iSCSI FE ports 94
MirrorView ports 94
MirrorView/A 147, 206
MirrorView/S 147, 148, 206
Mixing Fibre Channel configurations 101
Multi-Core Cache (MCC) 207
Multi-Core RAID (MCR) 208
multi-domain environments support 130
Multiple FCoE switch configurations 103

N

NDU 146
Network Data Management Protocol 139
Number of FC front-end ports 101
Number of FCoE front-end ports 101
Number of Flash (SSD) disks in FAST Cache 168, 195

O

operating systems supported 175
Optical cables 228

P

Path 96
Paths from a Server to an SP 113
Permanent Drive Sparing 209
Permit asynchronous login 131
permit asynchronous login 131
Physical memory per SP 197
Policies per storage system (max) 174, 201
Pool and pool LUN parameters 200
Pool LUN 120
Pool LUN size (max) 173, 200

Pool LUNs per storage system (max) 173, 200
 Pools 119
 Pools (max) 173, 200
 Port speed rules 98
 Portable Drives 209

R

R33 189
 RAID group and RAID group LUN (traditional)
 parameters 199
 RAID groups (max) 172, 199
 RAID types 173, 199, 200
 rail kits 214
 read/write cache allocation recommendations 167
 read/write cache configurations 165
 RecoverPoint write Splitter 186
 Replication and boot LUNs 136
 replication applications 154
 Replicator V2 Incremental Attach 138
 Reserved LUN pool 147
 Reserved LUNs per storage system 181
 retain domain security settings on exit 131

S

Sample VNX series configurations 118
 SAN Copy 150, 178, 201
 Scheduled events (max) 174, 201
 Single-initiator zoning 96
 Snapshots of snapshots 156
 Snapshots per source LUN 181, 204
 Snapshots per storage system 181
 SnapView configuration rules 154, 178, 201
 SnapView session source LUNs in consistent session
 181, 204
 SnapView sessions per source LUN 181, 204
 SnapView sessions per storage system 181, 204
 SnapView snapshot limits 181, 204
 Source LUNs per storage system 181, 204
 SP cache size per SP (max) 197
 SP write cache size per SP 198
 SPE 19
 SPE to DAE cables 230
 SPE to DAE SAS cables 230
 spin down 128
 Standby power supply (SPS) types 24
 storage domain 128

Storage group parameters 200
 Storage groups (max) 173, 200
 Storage system parameters for systems with VNX OE
 05.31 170, 196
 Storage tiering 124
 Storage-system initiator records required per Data
 Mover 98
 Storage-system port speed rules 98
 Supported FAST Cache RAID group and disk
 configurations 168, 195
 Symmetric Active/Active 209
 Symmetrix connections 88
 Symmetrix systems support 88

T

Thick LUNs 120
 Thin LUNs 121
 Thin Provisioning 169, 211
 True local rules 130
 TwinAx cables 229

U

Unified Configuration 13
 Unisphere 128
 Unisphere Quality of Service Manager 201

V

VAAI support for block 160
 VAAI support for NFS 157
 VASA support for VMware 159
 VG10 84
 VG2 80
 VG50 86
 VG8 82
 Virtual provisioning 119
 VNX Cable Types 225
 VNX configurations by model 30
 VNX data in place conversions 91
 VNX gateways 13
 VNX gateways and Symmetrix systems 88
 VNX hardware upgrades 89
 VNX models 12
 VNX RecoverPoint write Splitter usage 155
 VNX series MirrorView ports 94
 VNX Snapshot 155
 VNX Snapshots 205

Index

- VNX storage configuration options 13
- VNX Unified family 12
- VNX Unified hardware 12
- VNX5100 31
 - rules for FAST Cache 169
 - rules for Thin Provisioning 169
- VNX5300 36
- VNX5400 32, 40
- VNX5500 44
- VNX5600 49
- VNX5700 52
- VNX5800 57
- VNX7500 61
- VNX7600 69
- VNX8000 73
- vStorage API for Array Integration (VAI) 157, 160
- vStorage API for Storage Awareness (VASA) 159